Secure Cloud Storage with Data Dynamics and Privacy-Preserving Audits Using Secure Network Coding

Binanda Sengupta and Sushmita Ruj

Indian Statistical Institute, Kolkata, India
{binanda_r,sush}@isical.ac.in

arXiv:1612.08029v2 [cs.CR] 2 Jan 2017

Abstract. In the age of cloud computing, cloud users with a limited amount of storage can outsource their data to remote servers. The cloud servers, in lieu of monetary benefits, offer retrievability of their clients' data at any point of time. Secure cloud storage protocols ensure the integrity of the outsourced data that can be dynamic (or static) in nature depending on whether the client can (or cannot) update the uploaded data as needed. In this work, we explore the possibility of constructing a secure cloud storage for dynamic data by leveraging the idea of secure network coding. Specifically, we fail to provide a general construction of an *efficient* secure cloud storage protocol for dynamic data from an arbitrary secure network coding protocol. However, we show that some of the secure network coding schemes with some properties can be used to construct secure cloud storage protocols for dynamic data, and we indeed construct a publicly verifiable secure cloud storage protocol based on a secure network coding protocol. To the best of our knowledge, our scheme is the first secure cloud storage protocol for dynamic data that is based on a secure network coding protocol and that is secure in the standard model.

In a publicly verifiable setting, auditing task is often delegated to a third party auditor that audits the outsourced data on behalf of a client. It is desirable that the auditor gains no knowledge about the actual content of the client's data which may be sensitive. We extend our scheme in order to provide privacy-preserving audits where the content of the client's data is protected from the third party auditor. Furthermore, we extend our scheme in order to offer anonymity of a user updating shared data in an enterprise setting. In this setting, the cloud server cannot distinguish the user (belonging to a group) updating the data shared among the users of that group. We compare the performance of our secure cloud storage protocol with that of other secure cloud storage schemes and discuss some limitations of our scheme. Finally, we provide another construction of a secure cloud storage protocol that is specialized for append-only data and that overcomes some of the limitations of our earlier scheme.

Keywords: Cloud storage, network coding, dynamic data, shared data, public verifiability, privacy-preserving audits

1 Introduction

With the advent of cloud computing, cloud servers offer to their clients various facilities that include delegation of huge amount of computation and outsourcing large amount of data (say, in the order of terabytes). For example, a client (cloud user) having a smart phone with a low-performance processor or a limited amount of storage capacity cannot accomplish this heavy computation on its own or cannot store such a large volume of data in its own storage. The client can delegate her computation or storage to the cloud server. Now, the client only has to download the result of the computation or has to read (or update) the required portion of the uploaded data.

In case of storage outsourcing, the cloud server stores a massive volume of data on behalf of its clients (data owners). However, a malicious cloud server can delete some of the client's data (that are not accessed frequently) in order to save space. Secure cloud storage (SCS) protocols (two-party protocols between the client and the server) provide a mechanism to ensure that the client's data are stored untampered in the server. Depending on the nature of the data to be outsourced, secure cloud storage protocols are classified as: SCS protocols for *static* data (SSCS) [4,31,47] and SCS protocols for *dynamic* data (DSCS) [22,52,11,48]. For static data, the client cannot change her data after the initial outsourcing (suitable mostly for backup or archival data). Dynamic data are more generic in that the client can modify her data as often as needed. In SCS protocols, the client can audit her data stored in the server without accessing the whole data file, and still, be able to detect an unwanted modification of

the data done by a malicious server. The SCS protocols are *publicly verifiable* if the audits can be performed by any third party auditor (TPA) with the knowledge of public parameters only; they are *privately verifiable* if one needs some secret information of the client in order to perform an audit. In *privacy-preserving* audits (for publicly verifiable SCS protocols only), the TPA cannot learn the actual content of any portion of the data file.

In a network coding protocol [2,36], every intermediate node (all nodes except the source and target nodes) in a communication network combines the incoming packets to output another packet. These protocols enjoy much improved throughput, efficiency and scalability compared to the conventional store-and-forward routing where an incoming packet is relayed as it is. However, these protocols are prone to *pollution attacks* caused by malicious intermediate nodes that inject invalid packets in the network. These invalid packets produce more such packets downstream. In the worst case, the target node cannot decode the original file sent to it via the network. Secure network coding (SNC) protocols use cryptographic primitives in order to prevent these attacks. In an SNC protocol, the source node authenticates each of the packets to be transmitted through the network by attaching a small tag to it. These authentication tags are generated using homomorphic message authentication codes (MACs) [1] or homomorphic signatures [14,8,26,12]. Due to the homomorphic property of the tags, every intermediate node can combine the incoming packets to output another packet along with its authentication tag.

In a recent work, Chen et al. [16] show that, given a secure network coding (SNC) protocol, one can construct a secure cloud storage (SSCS) protocol for static data using the SNC protocol. However, their construction does not handle (more generic) dynamic data that makes it insufficient in many cloud applications where a client needs to update (insert, delete or modify) her data efficiently. Clearly, a naive way to update data in this scenario is to download the whole data file, perform the required updates and upload the file to the server again; but this procedure is highly inefficient as it requires huge amount of communication bandwidth for every update. Thus, further investigations are needed towards an efficient construction of a secure cloud storage (DSCS) protocol for dynamic data using an SNC protocol. In addition, there are various practical applications where *append-only* data need to be stored with a guarantee of retrievability (e.g., ledgers containing transactions, medical history of patients and different log data). Therefore, a more efficient solution (specifically for append-only data) would be helpful in this scenario.

Our Contribution In this paper, we provide a construction of a secure cloud storage protocol from a secure network coding (SNC) protocol. Our construction handles dynamic data, that is, the client can efficiently perform updates (insertion, deletion and modification) on her data outsourced to the cloud server. Our contributions are summarized as follows.

- We explore the possibility of providing a general construction of a DSCS protocol from any SNC protocol. We discuss about the challenges for such a general construction in details, and we identify some SNC protocols such that efficient DSCS protocols can be constructed using them. Indeed, we provide a basic construction of a DSCS protocol (DSCS I) from such an SNC protocol proposed by Catalano et al. [12]. Our construction is secure in the standard model and offers public verifiability.
- As our scheme is publicly verifiable, the client can delegate the auditing task to a third party auditor. She may also prefer her sensitive data not to be disclosed to the auditor. We extend our basic DSCS I construction in order to support privacy-preserving audits where the third party auditor cannot gain knowledge of the actual content of the data file owned by the client.
- In an enterprise setting, a group of users can share and update the outsourced data file. However, it is sometimes desirable to hide the identity of the user (from the cloud server) updating the data file. We extend DSCS I in order to preserve user-anonymity where the cloud server cannot distinguish the actual user in the group updating the file.
- We analyze the efficiency of DSCS I and compare it with other existing secure cloud storage protocols. We discuss about some limitations of an SNC-based SCS protocol (for static or dynamic data).
- We construct another publicly verifiable secure cloud storage protocol (DSCS II) for append-only data using an SNC protocol proposed by Boneh et al. [8]. This construction is secure in the random oracle model and overcomes some of the limitations of our earlier construction.

The rest of the paper is organized as follows. In Section 2, we discuss about the notations we use in this paper, and we briefly describe the secure network coding and the secure cloud storage protocols along with their

respective security models. Section 3 begins with a detailed discussion on the general construction of a DSCS protocol using an SNC protocol. Then, we describe our construction (DSCS I) that is based on a secure network coding protocol. In Section 4, we analyze the security of DSCS I and the probabilistic guarantees it offers. In Section 5, we extend our basic DSCS I scheme in order to support privacy-preserving audits. Section 6 describes a way in which DSCS I can be modified to preserve the user-anonymity in an enterprise setting. In Section 7, we analyze the efficiency of DSCS I and compare its performance with the existing secure cloud storage schemes. Section 8 describes our second construction of a secure cloud storage (DSCS II) that deals with append-only data. In the concluding Section 9, we summarize the work done in this paper.

2 Preliminaries and Background

2.1 Notation

We take λ to be the security parameter. An algorithm $\mathcal{A}(1^\lambda)$ is a probabilistic polynomial-time algorithm when its running time is polynomial in λ and its output y is a random variable which depends on the internal coin tosses of $\mathcal{A}$. An element a chosen uniformly at random from a set S is denoted as $a \xleftarrow{R} S$. A function $f : \mathbb{N} \to \mathbb{R}$ is called negligible in λ if for all positive integers c and for all sufficiently large λ, we have $f(\lambda) < \frac{1}{\lambda^c}$. In general, $\mathbb{F}$ is used to denote a finite field. The multiplication of a vector v by a scalar s is denoted by $s \cdot \mathrm{v}$. The terms *packet* and *vector* are used interchangeably in this work.

2.2 Secure Network Coding

Ahlswede et al. [2] introduce network coding as a replacement of the conventional store-and-forward routing for networks. In network coding, intermediate nodes (or routers) encode the received packets to output another packet which increases the throughput of the network (optimal in case of multicasting). Linear network coding was proposed by Li et al. [36]. Here, the file F to be transmitted is divided into several (say, m) packets (or vectors) $\mathrm{v}_1, \mathrm{v}_2, \ldots, \mathrm{v}_m$ each consisting of n components (or blocks), and each of these components is an element of a finite field $\mathbb{F}$. In other words, each $\mathrm{v}_i \in \mathbb{F}^n$ for $i \in [1, m]$. Then, the sender (or source) node augments each vector to form another vector $\mathrm{u}_i = [\mathrm{v}_i \ \mathrm{e}_i] \in \mathbb{F}^{n+m}$ for $i \in [1, m]$, where e_i is the m-dimensional unit vector containing 1 in the i-th position and 0 in others. Thus, the collection of m augmented vectors can be illustrated as

$$\begin{bmatrix} \text{—}\mathrm{u}_1\text{—} \\ \text{—}\mathrm{u}_2\text{—} \\ \vdots \\ \text{—}\mathrm{u}_m\text{—} \end{bmatrix} = \left[\begin{array}{c|c} \text{—}\mathrm{v}_1\text{—} & \text{—}\mathrm{e}_1\text{—} \\ \text{—}\mathrm{v}_2\text{—} & \text{—}\mathrm{e}_2\text{—} \\ \vdots & \vdots \\ \text{—}\mathrm{v}_m\text{—} & \text{—}\mathrm{e}_m\text{—} \end{array}\right] = \left[\begin{array}{cccc|cccc} v_{11} & v_{12} & \cdots & v_{1n} & 1 & 0 & \cdots & 0 \\ v_{21} & v_{22} & \cdots & v_{2n} & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ v_{m1} & v_{m2} & \cdots & v_{mn} & 0 & 0 & \cdots & 1 \end{array}\right].$$

Finally, the sender transmits these augmented vectors to the network. Let $V \subset \mathbb{F}^{n+m}$ be the linear subspace spanned by these augmented vectors $\mathrm{u}_1, \mathrm{u}_2, \ldots, \mathrm{u}_m$. A random file identifier `fid` is associated with the file F (or V). An intermediate node in the network outputs a linear combination of the received packets. In random (linear) network coding [29,28], the coefficients of these linear combinations are chosen uniformly at random from the field $\mathbb{F}$. An intermediate node in the network, upon receiving l packets $\mathrm{y}_1, \mathrm{y}_2, \ldots, \mathrm{y}_l \in \mathbb{F}^{n+m}$, chooses l coefficients $\nu_1, \nu_2, \ldots, \nu_l \xleftarrow{R} \mathbb{F}$ and outputs another packet $\mathrm{w} \in \mathbb{F}^{n+m}$ such that $\mathrm{w} = \sum_{i=1}^{l} \nu_i \cdot \mathrm{y}_i$ (here, summation refers to vector additions). Thus, the output packet of each intermediate node is of the form

$$\mathrm{w} = [w_1, w_2, \ldots, w_n, c_1, c_2, \ldots, c_m] \in V$$

for some $c_1, c_2, \ldots, c_m \in \mathbb{F}$, where $w_j = \sum_{i=1}^{m} c_i v_{ij}$ for each $j \in [1, n]$. When the receiver (or target) node accumulates m linearly independent vectors (or packets), it solves a system of linear equations to obtain the file destined to it.

A network coding protocol is vulnerable to pollution attacks, where a malicious intermediate node injects an invalid packet to the network and this invalid packet in turn pollutes multiple packets downstream. In a secure network coding (SNC) protocol, an authentication information (or tag) is attached to each packet. Every node can combine the incoming packets linearly to output another packet along with its authentication tag. Both the incoming packets and the output packet belong to the same linear subspace V. The authentication tags are computed using homomorphic message authentication codes (MACs) [1] or homomorphic signatures [14,8,26,6,12]. In an SNC protocol based on homomorphic signatures, every node in the network can verify the authenticity of each of its incoming vectors. On the other hand, in case of SNC protocols based on homomorphic MACs, it requires the knowledge of the secret key to verify an incoming vector. The property of aggregating (or combining) the tags for the incoming packets into a tag for the outgoing packet is required for the construction of a secure cloud storage protocol, where the cloud server has to send an aggregated tag for the response (see Section 3.3). So, we do not consider in our discussion SNC protocols based on (collision-resistant) homomorphic hashing, where intermediate nodes do not produce signatures for the outgoing packets (for example, [32,26]).

We define a secure network coding (SNC) protocol below. The values m, the dimension of the linear subspace $V \subset \mathbb{F}^{n+m}$ to sign, and n, the dimension of each vector v (before augmenting), are input to the procedures involved in the protocol.

Definition 1 (Secure Network Coding). *A secure network coding (SNC) protocol consists of the following algorithms.*

- SNC.KeyGen$(1^\lambda, m, n)$*: This procedure generates a secret key-public key pair* $K = (sk, pk)$ *for the sender.*
- SNC.TagGen$(V, sk, m, n, \texttt{fid})$*: On input a linear subspace* $V \subset \mathbb{F}^{n+m}$*, the secret key* sk *and a random file identifier* `fid` *associated with* V*, the sender runs this procedure to produce the authentication tag* t *for* V*.*
- SNC.Combine$(\{\mathrm{y}_i, t_i, \nu_i\}_{1 \leqslant i \leqslant l}, pk, m, n, \texttt{fid})$*: Given* l *incoming packets* $\mathrm{y}_1, \mathrm{y}_2, \ldots, \mathrm{y}_l \in \mathbb{F}^{n+m}$ *and their corresponding tags* $t_1, t_2, \ldots, t_l$ *for a file associated with* `fid`*, an intermediate node chooses* l *random coefficients* $\nu_1, \nu_2, \ldots, \nu_l \xleftarrow{R} \mathbb{F}$ *and runs this procedure. The procedure outputs another packet* $\mathrm{w} \in \mathbb{F}^{n+m}$ *and its authentication tag* t *such that* $\mathrm{w} = \sum_{i=1}^{l} \nu_i \cdot \mathrm{y}_i$*.*
- SNC.Verify$(\mathrm{w}, t, K, m, n, \texttt{fid})$*: An intermediate node or the receiver node, on input a packet* w *and its tag* t *for a file associated with* `fid`*, executes this procedure which returns 1 if* t *is authentic for the packet* w*; returns 0, otherwise.*

In some schemes, the procedure SNC.Verify requires only the public key pk [8,26,12]. The knowledge of the secret key sk is necessary to verify the incoming packets in other schemes [1].

Security of an SNC Protocol The security of an SNC protocol based on a homomorphic signature scheme is defined by the security game between a challenger and a probabilistic polynomial-time adversary $\mathcal{A}$ as stated below [12].

- **Setup** The adversary $\mathcal{A}$ provides the values m and n of its choice to the challenger. The challenger runs SNC.KeyGen$(1^\lambda, m, n)$ to output $K = (sk, pk)$ and returns pk to $\mathcal{A}$.
- **Queries** The adversary $\mathcal{A}$ specifies a sequence (adaptively chosen) of vector spaces $V_i \subset \mathbb{F}^{n+m}$ by respective augmented basis vectors $\{\mathrm{u}_{i1}, \mathrm{u}_{i2}, \ldots, \mathrm{u}_{im}\}$ and asks the challenger to authenticate the vector spaces. For each i, the challenger chooses a random file identifier $\texttt{fid}_i$ from a predefined space, generates an authentication tag t_i by running SNC.TagGen$(V_i, sk, m, n, \texttt{fid}_i)$ and gives t_i to $\mathcal{A}$.
- **Forgery** The adversary $\mathcal{A}$ outputs $(\texttt{fid}^*, \mathrm{w}^*, t^*)$.

Let the output vector $\mathrm{w}^* = [w_1^*, w_2^*, \ldots, w_{n+m}^*] \in \mathbb{F}^{n+m}$. Then, the adversary $\mathcal{A}$ wins the security game mentioned above if $[w_{n+1}^*, w_{n+2}^*, \ldots, w_{n+m}^*] \in \mathbb{F}^m$ is not the all-zero vector, SNC.Verify$(\mathrm{w}^*, t^*, pk, m, n, \texttt{fid}^*)$ outputs 1 and one of the following conditions is satisfied:
1. $\texttt{fid}^* \neq \texttt{fid}_i$ for all i (*type-1 forgery*)
2. $\texttt{fid}^* = \texttt{fid}_i$ for some i, but $\mathrm{w}^* \notin V_i$ (*type-2 forgery*).

For a secure network coding (SNC) protocol, the probability that the adversary $\mathcal{A}$ wins the security game is negligible in the security parameter λ.

We note that the security game for an SNC protocol based on homomorphic MACs is exactly the same as the game described above, except that the procedure SNC.KeyGen now produces a secret key only (unknown to $\mathcal{A}$) and the verification procedure SNC.Verify requires the knowledge of this secret key.

2.3 Secure Cloud Storage

Clients (cloud users) may want to outsource their huge amount of data to the cloud storage server. As the cloud service provider (possibly malicious) might discard old data to save some space, the clients need to be convinced that the outsourced data are stored untampered by the cloud server. A naive approach to ensure data integrity is that a client downloads the whole data from the server and verifies them individually segment by segment. However, this process is inefficient in terms of communication bandwidth required.

Building Blocks: PDP and POR Researchers come up with *proofs of storage* in order to resolve the issue mentioned above. Ateniese et al. [4] introduce the concept of *provable data possession* (PDP) where the client computes an authentication tag (for example, MAC) for each segment of her data (or file), and uploads the file along with the authentication tags. During an audit protocol, the client samples a predefined number of random segment-indices (challenge) and sends them to the server. We denote the cardinality of the challenge set by l which is typically taken to be $O(\lambda)$. The server does some computations (depending upon the challenge) over the stored data, and sends a proof (response) to the client who verifies the integrity of her data based on this proof. This scheme also introduces the notion of public verifiability[1] where the client (data owner) can delegate the auditing task to a third party auditor (TPA). Then, the TPA with the knowledge of the public key performs an audit. For privately verifiable schemes, only the client having knowledge of the secret key can verify the proof sent by the server. Other schemes achieving PDP include [5,22,52,51,23].

The first paper introducing *proofs of retrievability* (POR) for static data is by Juels and Kaliski [31] (a similar idea is given for sublinear authenticators by Naor and Rothblum [40]). According to Shacham and Waters [47], the underlying idea of a POR scheme is to encode the original file with an erasure code [37,44], authenticate the segments of the encoded file, and then upload them on the storage server. With this technique, the server has to delete or modify a considerable number of segments to actually delete or modify a data segment. This ensures that all segments of the file are retrievable from the responses of the server which passes an audit with some non-negligible probability. Following the work by Juels and Kaliski, several POR schemes have been proposed [10,20,49,11,48,13,3]. Some of these schemes are designed for static data, and the rest allow the client to modify data after the initial outsourcing.

As we deal with a single cloud server in this work, we only mention some secure cloud storage protocols in a distributed setting. Some of them include the works by Curtmola et al. [18] (using replication of data) and Bowers et al. [9] (using error-correcting codes and erasure codes). Dimakis et al. [19] introduce network coding in distributed storage systems where linear combinations of data segments are disseminated to multiple servers. In terms of repair bandwidth (bandwidth required to repair a failed server), this technique is more efficient than using the conventional erasure codes for distributing the segments. For such a distributed storage system, there are schemes for remote integrity checking [15,35,41] designed to achieve fast repair of a failed server.

We define an SCS protocol for dynamic data (DSCS) below [22]. A DSCS protocol can be privately verifiable if the verification algorithm of the protocol involves the secret key sk of the client; it is publicly verifiable, otherwise. In general, the term *verifier* is used to denote an auditor for a secure cloud storage. The client (for a privately verifiable protocol) or a third party auditor (for a publicly verifiable protocol) can act as the verifier.

[1] The term "public verifiability" discussed in this paper denotes (only) whether a third party auditor having the knowledge of the public parameters can perform an audit on behalf of the client (data owner). Following this notion of "public verifiability", there are schemes which are publicly verifiable [4,22,52,51,16]. We mention that this notion implicitly assumes that the client is honest. However, a malicious client can publish incorrect public parameters in order to get an honest server framed by a third party auditor [33].

Definition 2 (Secure Cloud Storage for Dynamic Data). *A secure cloud storage protocol for dynamic data (DSCS) consists of the following procedures.*

- DSCS.KeyGen(1^λ)*: This procedure generates a secret key-public key pair* $K = (sk, pk)$ *for the client.*
- DSCS.Outsource($F, K, m,$ `fid`)*: The client divides the file* F *associated with the file identifier* `fid` *into* m *segments and computes authentication tags for these segments using her secret key* sk*. Then, she constructs an authenticated data structure* M *on the authentication tags (for checking freshness of the data) and computes some metadata* d_M *for* M*. Finally, the client uploads the file* F' *(the file* F *and the authentication tags) along with* M *to the cloud storage server and stores* d_M *(and* m*) at her end.*
- DSCS.InitUpdate($i,$ `updtype`$, d_M, m,$ `fid`$)$*: The value of the variable* `updtype` *indicates whether the update is an insertion after or a modification of or the deletion of the* i*-th segment. Depending on the value of* `updtype`*, the client modifies* (m, d_M) *at her end and asks the server to perform the required update on the file associated with* `fid` *(related information specified in* `info`*).*
- DSCS.PerformUpdate($i,$ `updtype`$, F', M,$ `info`$, m,$ `fid`$)$*: The server performs the update on the file associated with* `fid` *and sends the client a proof* Π*.*
- DSCS.VerifyUpdate($i,$ `updtype`$, \Pi, m,$ `fid`$)$*: On receiving the proof* Π *for the file associated with* `fid` *from the server, the client checks whether* Π *is a valid proof.*
- DSCS.Challenge($pk, l,$ `fid`$)$*: During an audit, the verifier sends a challenge set* Q *of cardinality* $l = O(\lambda)$ *to the server.*
- DSCS.Prove($Q, pk, F', m,$ `fid`$)$*: The server, after receiving the challenge set* Q*, computes a proof of storage* T *corresponding to* Q *and a proof of freshness* Π*. Then, it sends* (T, Π) *to the verifier.*
- DSCS.Verify($Q, T, K, m,$ `fid`$)$*: The verifier checks if* T *is a valid proof of storage corresponding to the challenge set* Q *and* Π *is a valid proof of freshness. The verifier outputs 1 if both the proofs pass the verification; she outputs 0, otherwise.*

DSCS protocols based on PDP guarantee the extraction of *almost all* the segments of the file F. On the other hand, DSCS protocols based on POR ensure the extraction of *all* the segments of F with the help of erasure codes. Security of a DSCS protocol is discussed in Section 4.

3 Construction of an SCS Protocol for Dynamic Data Using an SNC Protocol

Chen et al. [16] propose a generic construction of a secure cloud storage protocol for static data from a secure network coding protocol. They consider the data file F to be stored in the server to be a collection of m vectors (or packets) each of which consists of n blocks. The underlying idea is to store these vectors (without augmenting them with unit vectors) along with their authentication tags in the server. During an audit, the client sends an l-element subset of the set of indices $\{1, 2, \ldots, m\}$ to the server. The server augments those vectors with the corresponding unit vectors, combines them linearly in an authenticated fashion and sends the output vector along with its tag to the client. Finally, the client verifies the authenticity of the received tag against the received vector. Thus, the server acts as an intermediate node, and the client acts as both the sender and the receiver (or the next intermediate router). We briefly discuss the algorithms involved in the general construction in Appendix A.

In a secure network coding protocol, the number of packets (or vectors) in the file to be transmitted through the network is fixed. This is because the length of the coefficient vectors used to augment the original vectors has to be determined a priori. That is why, a general construction of a secure cloud storage protocol as discussed above is suitable for static data in general. On the other hand, in a secure cloud storage protocol for dynamic data, clients can modify their data after they upload them to the cloud server initially. In this section, we discuss whether we can provide a general framework for constructing an efficient and secure cloud storage protocol for dynamic data (DSCS) from an SNC protocol.

3.1 On the General Construction of an Efficient DSCS Protocol from an SNC Protocol

In a secure network coding (SNC) protocol, a tag is associated with each packet such that the integrity of a packet can be verified using its tag. The SNC protocols found in the literature use homomorphic MACs [1] or

homomorphic signatures [14,8,26,6,12]. Following are the challenges in constructing an *efficient* DSCS protocol from these existing SNC protocols. We exclude, in our discussion, the work of Attrapadung and Libert [6] as their scheme is not efficient due to its reliance on (inefficient) composite-order bilinear groups.

1. *The DSCS protocol must handle the varying values of m appropriately*. In the network coding protocols mentioned above, the sender divides the file in m packets and augments them with unit coefficient vectors before sending them into the network. The length of these coefficient vectors is m which remains constant during transmission. In a secure cloud storage for dynamic data, the number of vectors may vary (for insertion and deletion). If we follow a similar general construction for a DSCS protocol as discussed above (for static data), we observe that the cloud server does not need to store the coefficient vectors. However, during an audit, the verifier selects a random l-element subset I of $[1, m]$ and the server augments the vectors with unit coefficient vectors of dimension m before generating the proof. Therefore, the verifier and the server need to keep an updated value of m.
   This issue can be resolved in the following way. The client includes the value of m in her public key and updates its value for each authenticated insertion or deletion. Thus, its latest value is known to the verifier and the server. We assume that, for consistency, the client (data owner) does not update her data during an audit.
2. *The index of a vector should not be embedded in its authentication tag*. In an SNC protocol, the file to be transmitted is divided into m packets $\mathrm{v}_1, \mathrm{v}_2, \ldots, \mathrm{v}_m$, where each $\mathrm{v}_i \in \mathbb{F}^n$ for $i \in [1, m]$ ([26] replaces $\mathbb{F}$ by $\mathbb{Z}$). The sender augments each vector to form another vector $\mathrm{u}_i = [\mathrm{v}_i\ \mathrm{e}_i] \in \mathbb{F}^{n+m}$ for $i \in [1, m]$, where e_i is the m-dimensional unit vector containing 1 in i-th position and 0 in others. Let $V \subset \mathbb{F}^{n+m}$ be the linear subspace spanned by these augmented basis vectors $\mathrm{u}_1, \mathrm{u}_2, \ldots, \mathrm{u}_m$. The sender authenticates the subspace V by authenticating these augmented vectors before transmitting them to the network [14,1,8,26,12]. In a scheme based on homomorphic MACs [1], the sender generates a MAC for the i-th basis vector u_i and the index i serves as an input to the MAC algorithm (for example, i is an input to the pseudorandom function in [1]). On the other hand, for the schemes based on homomorphic signatures, the sender generates a signature t_i on the i-th basis vector u_i. In some schemes based on homomorphic signatures, the index i is embedded in the signature t_i on the i-th augmented vector. For example, $H(\mathtt{fid}, i)$ is embedded in t_i [8,26], where $\mathtt{fid}$ is the file identifier and H is a hash function modeled as a random oracle. Section 8.1 gives a brief overview of the scheme proposed by Boneh et al. [8] and shows how each t_i includes the value $H(\mathtt{fid}, i)$ (see Eqn. 10).
   These schemes are not suitable for the construction of an *efficient* DSCS protocol due to the following reason. For dynamic data, the client can insert a vector in a specified position or delete an existing vector from a specified location. In both cases, the indices of the subsequent vectors are changed. Therefore, the client has to download all these subsequent vectors and compute fresh authentication tags for them before uploading the new vector-tag pairs to the cloud server. This makes the DSCS protocol inefficient. However, in a few schemes, instead of hashing vector indices as in [8,26], there is a one-to-one mapping from the set of indices to some group [14,12], and these group elements are made public. This increases the size of the public key of these schemes. However, an efficient DSCS protocol can be constructed from them. In fact, we construct a DSCS protocol (described in Section 3.3) based on the SNC protocol proposed by Catalano et al. [12]. We note that Chen et al. [16] construct an SCS protocol from the same SNC protocol, but for static data only.
3. *The freshness of data must be guaranteed*. The freshness of storage requires that the server is storing an up-to-date version of the data file. For dynamic data, the client can modify an existing vector. However, a malicious cloud server may discard this change and keep an old copy of the vector. As the old copy of the vector and its corresponding tag are valid, the client has no way to detect if the cloud server is storing the latest copy.
   We ensure the freshness of the client's data, in our DSCS construction, using an *authenticated data structure (rank-based authenticated skip list) on the authentication tags of all the vectors*. In other words, the authenticity of the vectors is maintained by their tags, and the integrity of the tags is in turn maintained by the skip list. The advantage of building the skip list on the tags (over building it on the vectors) is that the tags are much shorter than a vector, and this decreases the size of the proof sent by the server. When a vector is inserted (or modified), its tag is also updated and sent to the server. The server updates the skip list accordingly. For deletion of a vector, the server simply removes the corresponding tag from the skip list. Finally, the server sends to the client a proof of performing the required update properly. We briefly discuss, in Section 3.2, about rank-based authenticated skip lists that we use in our construction described in Section 3.3.

In addition to the requirements mentioned above, it is often desired that a DSCS protocol (an SCS protocol, in general) satisfies the following properties.

4. *Public verifiability* For a publicly verifiable DSCS protocol, the auditing task can be delegated to a third party auditor (TPA). In a secure network coding protocol built on homomorphic MACs, the secret key (for example, the secret keys of the pseudorandom generator and the pseudorandom function in [1]) is needed to verify the authenticity of an incoming packet. This property restricts the secure cloud storage protocol built on such an SNC protocol to be privately verifiable only.
5. *Privacy-preserving audits* In privacy-preserving audits (for a publicly verifiable DSCS protocol), the TPA cannot gain the knowledge of the challenged vectors.

3.2 Rank-Based Authenticated Skip Lists

For dynamic data, we need some authenticated data structures like Merkle hash trees [38], rank-based authenticated skip lists [22] and rank-based RSA trees [42,22] to verify the freshness of each of the vectors. Erway et al. [22] propose rank-based authenticated skip lists based on labeled skip lists [27,43]. We give a brief introduction to the procedures of rank-based authenticated skip lists stored remotely in a server as follows.

– ListInit$(t_1, \ldots, t_m)$: Let $\{t_1, \ldots, t_m\}$ be an ordered list of m elements on which a rank-based authenticated skip list M is to be built. These elements are kept in the bottom-level nodes of the skip list in an ordered fashion. For each node z of the skip list: $right(z)$ and $down(z)$ are two pointers to the successors of z, $rank(z)$ is the number of bottom-level nodes reachable from z (including z if z itself is a bottom-level node), $high(z)$ and $low(z)$ are the indices of the leftmost and rightmost bottom-level nodes reachable from z, $f(z)$ is the label associated with the node z, and $l(z)$ is the level of z ($l(z) = 0$ for a bottom-level node z).
Initially, all these information (except the label) are computed for each node in the skip list. In addition, the i-th bottom-level node z contains $x(z) = t_i, \forall i \in [1, m]$. Finally, for each node z, the label $f(z)$ is computed using a *collision-resistant* hash function h as

$$f(z) = \begin{cases} 0, & \text{if } z \text{ is } \texttt{null} \\ f_1, & \text{if } l(z) = 0 \\ f_2, & \text{if } l(z) > 0 \end{cases}$$

where $f_1 = h(l(z)||rank(z)||x(z)||f(right(z)))$ and $f_2 = h(l(z)||rank(z)||f(down(z))||f(right(z)))$. Fig. 1 illustrates a rank-based authenticated skip list for an ordered list $\{t_1, \ldots, t_9\}$.
The skip list along with all the associated information are stored in the server. The client only stores the value of m and the label of the root node r (i.e., $f(r)$) as the metadata d_M.
– ListAuthRead(i, m): When the client wants to read the i-th element t_i, the server sends the requested element along with a proof $\Pi(i)$ to the client. Let the verification path of the i-th element be a sequence of nodes $z_1, \ldots, z_k$, where z_1 is the bottom-level node storing the i-th element and $z_k = r$ is the root node of the skip list (see Fig. 1). Then, the proof $\Pi(i)$ is of the form

$$\Pi(i) = (A(z_1), \ldots, A(z_k)),$$

where $A(z) = (l(z), q(z), d(z), g(z))$. Here, $l(z)$ is the level of the node z, $d(z)$ is 0 (or 1) if $down(z)$ (or $right(z)$) points to the previous node of z in the sequence, and $q(z)$ and $g(z)$ are the rank and label (respectively) of the successor node of z that is not present on the verification path.
– ListVerifyRead$(i, d_M, t_i, \Pi(i), m)$: Upon receiving the proof $(t_i, \Pi(i))$ from the server, the client checks if the proof corresponds to the latest metadata d_M stored at her end. The client outputs 1 if the proof matches with the metadata; she outputs 0, otherwise.
Due to the collision-resistance property of the hash function h that is used to generate the labels of the nodes of the skip list, the server cannot pass the verification without storing the element t_i properly, except with some probability negligible in the security parameter λ.

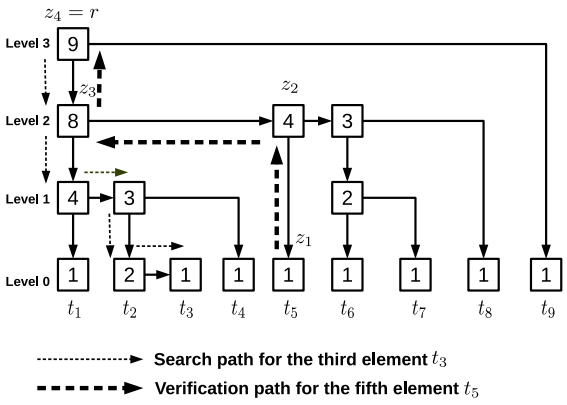

Fig. 1. The structure of a rank-based authenticated skip list. The skip list is built on an ordered list $\{t_1, \ldots, t_9\}$. The root node of the skip list is r. The $rank$ of each node in the list is written inside it. The elements are in the bottom-level (Level 0) nodes, and the root node r resides in the highest level (Level 3). The search path for the third element t_3 and the verification path for the fifth element t_5 are shown.

- ListInitUpdate$(i, \texttt{updtype}, d_M, t'_i, m)$: An update can be an insertion after or a modification of or the deletion of the i-th bottom-level node. The type of the update is stored in a variable `updtype`. The client defines $j = i$ (for an insertion or modification) or $j = i - 1$ (for a deletion). She calls ListAuthRead(j, m) for the existing skip list M and verifies the response sent by the server by calling ListVerifyRead$(j, d_M, t_j, \Pi(j), m)$. If the proof does not match with the metadata d_M (the label of the root node of the existing skip list M), she aborts. Otherwise, she updates the value of m, computes the new metadata d'_M using the proof and stores it at her end temporarily. Then, the client asks the server to perform the update specifying the location i, `updtype` (insertion, deletion or modification) and the new element t'_i (`null` for deletion).
- ListPerformUpdate$(i, \texttt{updtype}, t'_i, M)$: Depending on the value of `updtype`, the server performs the update asked by the client, computes a proof Π similar to the one generated during ListAuthRead and sends Π to the client.
- ListVerifyUpdate$(i, \texttt{updtype}, t'_i, d'_M, \Pi, m)$: On receiving the proofs from the server, the client verifies the proof Π and computes the new metadata d_{new} based on Π. If $d'_M = d_{new}$ and Π is a valid proof, the client sets $d_M = d'_M$, deletes the temporary value d'_M and outputs 1. Otherwise, she changes m to its previous value, deletes d'_M and outputs 0.

Note One limitation of using dynamic versions of Merkle hash trees (for example, authenticated red-black trees) instead of rank-based authenticated skip lists is that a series of insertions after a particular location makes the tree imbalanced and increases the height of the tree by the number of insertions. In the two-party model as in our case, no efficient rebalancing techniques (for updating the authentication information of the affected nodes efficiently) for such a tree have been studied [22]. On the other hand, due to the properties of a skip list [43], the number of levels in a skip list is logarithmic in m with high probability. For this reason, the size of a proof, the computation time for the server and the verification time for the client are $O(\log m)$ with high probability.

3.3 DSCS I: A DSCS Protocol Using an SNC Protocol

Following the work of Chen et al. [16], we construct a secure cloud storage protocol for dynamic data (DSCS I) from the secure network coding (SNC) protocol proposed by Catalano et al. [12] which is secure in the standard model. This basic construction exploits a rank-based authenticated skip list (discussed in Section 3.2) to ensure the freshness of the dynamic data. DSCS I consists of the following procedures. Let h be the *collision-resistant* hash function used in the rank-based authenticated skip list we use in our construction. We assume that the file F to be outsourced to the server is a collection of m vectors where each of the vectors consists of n blocks. We note that

the procedures KeyGen, Outsource, Prove and Verify in DSCS I call the procedures SNC.KeyGen, SNC.TagGen, SNC.Combine and SNC.Verify (respectively) of the underlying SNC protocol [12] along with performing some other operations related to the authenticated data structure. Here, we provide a detailed description.

- KeyGen($1^\lambda, m, n$): The client selects two random safe primes[2] p, q of length $\lambda/2$ bits each and takes $N = pq$. The client chooses another random prime e of length $\lambda + 1$ (in bits) and sets the file identifier `fid` to be equal to e. She selects $g, g_1, \ldots, g_n, h_1, \ldots, h_m \xleftarrow{R} \mathbb{Z}_N^*$. The secret key sk is (p, q), and the public key pk consists of $(N, e, g, g_1, \ldots, g_n, h_1, \ldots, h_m, d_M, m, n)$. Initially, d_M is `null`. Let $K = (sk, pk)$.
- Outsource($F, K,$ `fid`): The file F (associated with the identifier `fid`) consists of m vectors each of them having n blocks. We assume that each of these blocks is an element of $\mathbb{F}_e$. Then, for each $1 \leqslant i \leqslant m$, the i-th vector v_i is of the form $[v_{i1}, \ldots, v_{in}] \in \mathbb{F}_e^n$. For each vector v_i, the client selects a random element $s_i \xleftarrow{R} \mathbb{F}_e$ and computes x_i such that

$$x_i^e = g^{s_i}(\prod_{j=1}^{n} g_j^{v_{ij}})h_i \bmod N. \tag{1}$$

  Now, $t_i = (s_i, x_i)$ acts as an authentication tag for the vector v_i. The client constructs a rank-based authenticated skip list M on the authentication tags $\{t_i\}_{1 \leqslant i \leqslant m}$ and computes the metadata d_M (the label of the root node of M). Finally, the client updates d_M in the public key and uploads the file $F' = \{(\mathrm{v}_i, t_i)\}_{1 \leqslant i \leqslant m}$ along with M to the cloud server.
- InitUpdate($i,$ `updtype`$, pk,$ `fid`): The value of the variable `updtype` indicates whether the update is an insertion after or a modification of or the deletion of the i-th vector. The client performs one of the following operations depending on the value of `updtype`.
  1. If `updtype` is insertion, the client selects $h' \xleftarrow{R} \mathbb{Z}_N^*$ and generates the new vector-tag pair (v', t'). She runs ListInitUpdate on $(i,$ `updtype`$, d_M, t', m)$ and sends (h', v') to the server.
  2. If `updtype` is modification, the client generates the new vector-tag pair (v', t'). Then, the client runs ListInitUpdate($i,$ `updtype`$, d_M, t', m$) and sends v' to the server.
  3. If `updtype` is deletion, the client runs the procedure ListInitUpdate on $(i,$ `updtype`$, d_M, t', m)$, where t' is `null`.

  The client stores the value of the new metadata d'_M temporarily at her end.
- PerformUpdate($i,$ `updtype`$, F', M, h', \mathrm{v}', t', pk,$ `fid`): We assume that, for efficiency, the server keeps a local copy of the ordered list of h_j values for $1 \leqslant j \leqslant m$. Based on the value of `updtype`, the server performs one of the following operations.
  1. If `updtype` is insertion, the server sets $m = m + 1$, inserts h' in the $(i + 1)$-th position in the list of h_j values (for $1 \leqslant j \leqslant m$) and inserts v' after the i-th vector. The server runs ListPerformUpdate on the input $(i,$ `updtype`$, t', M)$.
  2. If `updtype` is modification (h' is `null`), the server modifies the i-th vector to v' and runs the procedure ListPerformUpdate on $(i,$ `updtype`$, t', M)$.
  3. If `updtype` is deletion (h', v' and t' are `null`), the server sets $m = m - 1$, deletes the particular h_i value from the list of h_j values ($j \in [1, m]$) and runs ListPerformUpdate($i,$ `updtype`$,$ `null`$, M$).
- VerifyUpdate($i,$ `updtype`$, t', d'_M, \Pi, pk,$ `fid`): After receiving the proof from the server, the client performs ListVerifyUpdate($i,$ `updtype`$, t', d'_M, \Pi, m$). If the output of ListVerifyUpdate is 1, the client outputs 1 and updates her public key (the latest values of m, d_M and h_j for $j \in [1, m]$) accordingly. Otherwise, the client outputs 0.
- Challenge($pk, l,$ `fid`): During an audit, the verifier selects I, a random l-element subset of $[1, m]$. Then, she generates a challenge set $Q = \{(i, \nu_i)\}_{i \in I}$, where each $\nu_i \xleftarrow{R} \mathbb{F}_e$. The verifier sends the challenge set Q to the cloud server.
- Prove($Q, pk, F', M,$ `fid`): The cloud server, after receiving the challenge set $Q = \{(i, \nu_i)\}_{i \in I}$, computes $s = \sum_{i \in I} \nu_i s_i \bmod e$ and $s' = (\sum_{i \in I} \nu_i s_i - s)/e$. The server, for each $i \in I$, forms $\mathrm{u}_i = [\mathrm{v}_i\ \mathrm{e}_i] \in \mathbb{F}_e^{n+m}$

[2] A safe prime is a prime of the form $2p' + 1$, where p' is also a prime.

by augmenting the vector v_i with the unit coefficient vector e_i. Then, it computes $\mathrm{w} = \sum_{i \in I} \nu_i \cdot \mathrm{u}_i \bmod e \in \mathbb{F}_e^{n+m}$, $\mathrm{w}' = (\sum_{i \in I} \nu_i \cdot \mathrm{u}_i - \mathrm{w})/e \in \mathbb{F}_e^{n+m}$ and

$$x = \frac{\prod_{i \in I} x_i^{\nu_i}}{g^{s'} \prod_{j=1}^{n} g_j^{w'_j} \prod_{j=1}^{m} h_j^{w'_{n+j}}} \bmod N. \tag{2}$$

Let $\mathrm{y} \in \mathbb{F}_e^n$ be the first n entries of w and $t = (s, x)$. The server sends $T = (T_1, T_2)$ as a proof of storage corresponding to the challenge set Q, where $T_1 = (\mathrm{y}, t)$ and $T_2 = \{(t_i, \Pi(i))\}_{i \in I}$.

– Verify$(Q, T, pk, \mathtt{fid})$: Using $Q = \{(i, \nu_i)\}_{i \in I}$ and $T = (\mathrm{y}, t)$ sent by the server, the verifier constructs a vector $\mathrm{w} = [w_1, \ldots, w_n, w_{n+1}, \ldots, w_{n+m}] \in \mathbb{F}_e^{n+m}$, where the first n entries of w are the same as those of y and the $(n+i)$-th entry is ν_i if $i \in I$ (0 if $i \notin I$). Then, the verifier checks whether

$$x^e \stackrel{?}{=} g^s \prod_{j=1}^{n} g_j^{w_j} \prod_{j=1}^{m} h_j^{w_{n+j}} \bmod N. \tag{3}$$

She also verifies if, for each $i \in I$, $\Pi(i)$ is a valid proof (with respect to d_M) for t_i. The verifier outputs 1 if the proof passes all the verifications; she outputs 0, otherwise.

The DSCS protocol described above is publicly verifiable as only the knowledge of the public key of the client (data owner) enables one to perform an audit (see the footnote in Section 2.3).

4 Security of DSCS I

4.1 Overview of Security of a DSCS Protocol

A secure DSCS protocol must satisfy the following properties [22,48]. A formal security model is described in Section 4.2.

1. **Authenticity** The authenticity of storage requires that the cloud server cannot produce a valid proof of storage T' (corresponding to the challenge set Q) without storing the challenged segments and their respective authentication information untampered, except with a probability negligible in λ.
2. **Freshness** The freshness of storage guarantees that the server is storing an up-to-date version of the file F.
3. **Extractability** The extractability (or retrievability) of data requires that, given a probabilistic polynomial-time adversary $\mathcal{A}$ that can respond correctly to a challenge Q with some non-negligible probability, there exists a polynomial-time extractor algorithm $\mathcal{E}$ that can extract (at least) the challenged segments (except with negligible probability) by challenging $\mathcal{A}$ for a polynomial (in λ) number of times and verifying the responses sent by $\mathcal{A}$. The algorithm $\mathcal{E}$ has a non-black-box access to $\mathcal{A}$. Thus, $\mathcal{E}$ can rewind $\mathcal{A}$, if required.

4.2 Security Model

The DSCS I protocol offers the guarantee of dynamic provable data possession (DPDP) [22]. We describe the data possession game of DPDP between the challenger (acting as the client) and the adversary (acting as the cloud server) as follows.

– The challenger generates a key pair (sk, pk) and gives pk to the adversary.
– The adversary selects a file F associated with the identifier $\mathtt{fid}$ to store. The challenger processes the file to form another file F' with the help of sk and returns F' to the adversary. The challenger stores only some metadata to verify the updates to be performed by the adversary later. The adversary chooses a sequence of updates (of its choice) defined by $(\mathtt{updtype}_i, \mathtt{info}_i)$ for $1 \leqslant i \leqslant q_1$ (q_1 is polynomial in the security parameter λ) and asks the challenger to initiate the update. For each update, the challenger runs InitUpdate and stores the latest metadata at her end. The adversary sends a proof after executing PerformUpdate. The challenger verifies this proof by running VerifyUpdate and updates her metadata if and only if the proof passes the verification. The adversary is notified about the output of VerifyUpdate for each update.

– Let F^* be the final state of the file after q_1 updates. The challenger has the latest metadata for the file F^*. Now, she challenges the adversary with a random challenge set Q, and the adversary returns a proof $T = (T_1, T_2)$ to the challenger. The adversary wins the game if the proof passes the verification. The challenger can challenge the adversary q_2 (polynomial in λ) number of times in an attempt to extract (at least) the challenged vectors of F^*.

Definition 3 (Security of Dynamic Provable Data Possession). *A dynamic provable data possession scheme is secure if, given any probabilistic polynomial-time adversary $\mathcal{A}$ who can win the data possession game mentioned above with some non-negligible probability, there exists a polynomial-time extractor algorithm $\mathcal{E}$ that can extract (at least) the challenged vectors of the file by interacting (via challenge-response) with $\mathcal{A}$ polynomially many times.*

4.3 Security Analysis

We state and prove the following theorem in order to analyze the security of DSCS I.

Theorem 1. *Given that the hash function used to construct the rank-based authenticated skip list is collision-resistant and the underlying network coding scheme is secure, the DSCS I protocol described in Section 3.3 is secure in the standard model according to Definition 3.*

Proof. We use the following claims in order to prove Theorem 1.

Claim. Given that the hash function used to construct the rank-based authenticated skip list is collision-resistant, the freshness of the data file is guaranteed in the DSCS I protocol.

Proof. The freshness of storage requires that the cloud server must store an up-to-date version of the data file outsourced by the client. In DSCS I, the hash function h (see Section 3.2) used to compute the labels of the nodes in the rank-based authenticated skip list M is collision-resistant. For each update, freshness of data is guaranteed using the procedure VerifyUpdate (by computing d_{new} from Π and checking if $d'_M \stackrel{?}{=} d_{new}$) during the data possession game. Moreover, for each challenge Q, freshness of data is guaranteed by checking the validity of the proof $T_2 = \{(t_i, \Pi(i))\}_{i \in I}$ with respect to the latest metadata d_M (in the procedure Verify) during the data possession game and the extraction phase. If the adversary is able to forge a skip list proof with respect to the latest value of d_M (i.e., if it makes the output of some execution of VerifyUpdate or Verify to be 1 without storing the latest authentication tags), then it must have found a collision for the hash function h in some level of the rank-based authenticated skip list. This event occurs with a negligible probability as h is taken to be collision-resistant. □

Claim. Given that the underlying network coding scheme is secure in the standard model, the authenticity of the data file is guaranteed in DSCS I.

Proof. The authenticity of storage demands that the cloud server, without storing the challenged vectors and their respective authentication tags appropriately, cannot produce a valid response $T' = (T'_1, T'_2) = ((\mathrm{y}', t'), T'_2)$ for a challenge set $Q = \{(i, \nu_i)\}_{i \in I}$ during the data possession game (and during the extraction phase). The data file F with a random (but unique) $\mathtt{fid}$ is identified by the augmented vectors $\mathrm{u}_i = [\mathrm{v}_i \ \mathrm{e}_i] \in \mathbb{F}_e^{n+m}$ for $i \in [1, m]$. Let $T = (T_1, T_2) = ((\mathrm{y}, t), T_2)$ be the response computed honestly; thus, Verify$(Q, T, pk, \mathtt{fid}) = 1$. Let $\mathrm{w} = [w_1, \ldots, w_n, w_{n+1}, \ldots, w_{n+m}] \in \mathbb{F}_e^{n+m}$, where the first n entries of w are the same as those of y and the $(n+i)$-th entry is ν_i if $i \in I$ (0 if $i \notin I$).

Now, if possible, we assume that the adversary produces another valid response $T' = (T'_1, T'_2) = ((\mathrm{y}', t'), T'_2)$ for the same Q such that Verify$(Q, T', pk, \mathtt{fid}) = 1$, where $\mathrm{y} \neq \mathrm{y}'$. Let $\mathrm{w}' = [w'_1, \ldots, w'_n, w'_{n+1}, \ldots, w'_{n+m}] \in \mathbb{F}_e^{n+m}$, where the first n entries of w′ are the same as those of y′ and the $(n+i)$-th entry is ν_i if $i \in I$ (0 if $i \notin I$). Clearly, $\mathrm{w} \neq \mathrm{w}'$. We observe that the procedure Verify executes the procedure SNC.Verify. As Verify$(Q, T', pk, \mathtt{fid})$ outputs 1, it follows that SNC.Verify$(\mathrm{w}', t', pk, m, n, \mathtt{fid}) = 1$.

We also note that, for a particular challenge Q, the set of indices I and the corresponding coefficients ν_i (for $i \in I$) are randomly chosen by the challenger (data possession game) or by the extractor (extraction phase). As

the basis vectors $\mathrm{u}_1, \ldots, \mathrm{u}_m$ for the particular data file F (identified by `fid`) are unique, their linear combination using fixed coefficients (i-th coefficient is ν_i or 0 depending on whether $i \in I$ or $i \notin I$) is also unique. This unique linear combination is w. Therefore, $\mathrm{w}' \notin \mathrm{span}(\mathrm{u}_1, \ldots, \mathrm{u}_m)$.

To sum up, we find a pair (w', t') for a file with identifier `fid` such that $[w'_{n+1}, w'_{n+2}, \ldots, w'_{n+m}] \in \mathbb{F}_e^m$ is not the all-zero vector, SNC.Verify$(\mathrm{w}', t', pk, m, n,$ `fid`$)$ outputs 1 and $\mathrm{w}' \notin \mathrm{span}(\mathrm{u}_1, \ldots, \mathrm{u}_m)$ (*type-2 forgery* as mentioned in Section 2.2). However, since the network coding protocol [12] we use in DSCS I is secure in the standard model, the adversary cannot produce such a response $T'_1 = (\mathrm{y}', t')$, except with some probability negligible in the security parameter. □

We define a polynomial-time extractor algorithm $\mathcal{E}$ that can extract (at least) the challenged vectors (except with negligible probability) by interacting with an adversary $\mathcal{A}$ that wins the data possession game mentioned above with some non-negligible probability. As DSCS I satisfies the *authenticity* and *freshness* properties mentioned above, $\mathcal{A}$ cannot produce a proof $T = (T_1, T_2)$ for a given challenge set $Q = \{(i, \nu_i)\}_{i \in I}$ without storing the challenged blocks and their corresponding tags properly, except with some negligible probability. This means that if the output of the procedure Verify is 1 during the extraction phase, the vector y in the proof is the linear combination of the original data vectors v_i for $i \in I$ using coefficients $\{\nu_i\}_{i \in I}$.

Suppose that the extractor $\mathcal{E}$ wants to extract l blocks indexed by J. It challenges $\mathcal{A}$ with $Q = \{(i, \nu_i)\}_{i \in J}$. If the proof is valid (checked using Verify), $\mathcal{E}$ initializes a matrix $M_{\mathcal{E}}$ as $[\nu_{1i}]_{i \in J}$, where $\nu_{1i} = \nu_i$ for each $i \in J$. The extractor challenges $\mathcal{A}$ for the same J but with different random coefficients. If the procedure Verify outputs 1 and the vector of coefficients is linearly independent to the existing rows of $M_{\mathcal{E}}$, then $\mathcal{E}$ appends this vector to $M_{\mathcal{E}}$ as a row. The extractor $\mathcal{E}$ runs this procedure until the matrix $M_{\mathcal{E}}$ has l linearly independent rows. So, the final form of the full-rank matrix $M_{\mathcal{E}}$ is $[\nu_{ji}]_{j \in [1,l], i \in J}$. Therefore, the challenged blocks can be extracted with the help of Gaussian elimination. This completes the proof of Theorem 1. ■

4.4 Probabilistic Guarantees

If the server has corrupted a constant fraction (say, β) of vectors in a file, then the server passes an audit with probability $p_{cheat} = (1 - \beta)^l$, where l is the cardinality of Q. The probability p_{cheat} is very small for large values of l. Typically, l is taken to be $O(\lambda)$ in order to make the probability p_{cheat} negligible in λ. Thus, the verifier detects a malicious server corrupting β-fraction of the file with probability $p_{detect} = 1 - p_{cheat} = 1 - (1 - \beta)^l$, and it guarantees the integrity of *almost all* vectors of the file.

5 Enabling Privacy-Preserving Audits in DSCS I

The basic DSCS I protocol described in Section 3.3 is publicly verifiable, that is, a third party auditor (TPA) having the knowledge of the public parameters can perform an audit. Chen et al. [16] construct a secure cloud storage protocol for *static* data using the same SNC protocol [12]. They note that, in order to make an audit privacy-preserving, the server adds a random linear combination of some random vectors to the computed value of y to form the final response. Due to the addition of this random component to the resulting vector y, the third party auditor (TPA) cannot gain knowledge (by solving a set of linear equations) of the challenged vectors. Accordingly, we modify our basic DSCS I protocol in order to support privacy-preserving audits in the following way. We change two procedures, Outsource and Prove, from the basic protocol. The rest of the procedures involved in the protocol are the same as in the basic protocol.

– Outsource$(F, K,$ `fid`$)$: The file F (associated with the identifier `fid`) consists of m vectors each of them having n blocks. We assume that each of these blocks is an element of $\mathbb{F}_e$. Then, for each $1 \leqslant i \leqslant m$, the i-th vector v_i is of the form $[v_{i1}, \ldots, v_{in}] \in \mathbb{F}_e^n$. For each vector v_i, the client selects a random element $s_i \xleftarrow{R} \mathbb{F}_e$ and computes x_i such that

$$x_i^e = g^{s_i} (\prod_{j=1}^{n} g_j^{v_{ij}}) h_i \bmod N. \tag{4}$$

Now, $t_i = (s_i, x_i)$ acts as an authentication tag for the vector v_i. The client constructs a rank-based authenticated skip list M on the authentication tags $\{t_i\}_{1 \leqslant i \leqslant m}$ and computes the metadata d_M (the label of the root node of M). Finally, the client updates d_M in the public key and uploads the file $F' = \{(\mathrm{v}_i, t_i)\}_{1 \leqslant i \leqslant m}$ along with M to the cloud server. The client also uploads, to the server, a set $\tilde{I} = \{\tilde{\mathrm{u}}_1, \ldots, \tilde{\mathrm{u}}_k\}$ of k vectors (along with their authentication tags) generated as follows. For each $i \in [1, k]$, $\tilde{\mathrm{u}}_i = [\tilde{u}_{i1}, \ldots, \tilde{u}_{i(n+m)}]$, where $\tilde{u}_{ij} \xleftarrow{R} \mathbb{F}_e$ for each $j \in [1, n]$ and $\tilde{u}_{ij} = 0$ for each $j \in [n+1, n+m]$. For each vector $\tilde{\mathrm{u}}_i$, the client selects a random element $\tilde{s}_i \xleftarrow{R} \mathbb{F}_e$ and computes $\tilde{x}_i$ such that

$$\tilde{x}_i^e = g^{\tilde{s}_i} \prod_{j=1}^{n} g_j^{\tilde{u}_{ij}} \bmod N. \tag{5}$$

Then, the client uploads the set $\tilde{I}$ along with the tags $\{(\tilde{s}_i, \tilde{x}_i)\}_{i \in [1,k]}$ to the cloud server.

– Prove$(Q, pk, F', M, \mathtt{fid})$: The cloud server, after receiving the challenge set $Q = \{(i, \nu_i)\}_{i \in I}$, selects a set $\tilde{Q} = \{\beta_1, \ldots, \beta_k\}$, where $\beta_i \xleftarrow{R} \mathbb{F}_e$ for each $1 \leqslant i \leqslant k$. Then, it computes $s = (\sum_{i \in I} \nu_i s_i + \sum_{1 \leqslant i \leqslant k} \beta_i \tilde{s}_i) \bmod e$ and $s' = (\sum_{i \in I} \nu_i s_i + \sum_{1 \leqslant i \leqslant k} \beta_i \tilde{s}_i - s)/e$. The server, for each $i \in I$, forms $\mathrm{u}_i = [\mathrm{v}_i \ \mathrm{e}_i] \in \mathbb{F}_e^{n+m}$ by augmenting the vector v_i with the unit coefficient vector e_i. Then, it computes $\mathrm{w} = (\sum_{i \in I} \nu_i \cdot \mathrm{u}_i + \sum_{1 \leqslant i \leqslant k} \beta_i \cdot \tilde{\mathrm{u}}_i) \bmod e \in \mathbb{F}_e^{n+m}$, $\mathrm{w}' = (\sum_{i \in I} \nu_i \cdot \mathrm{u}_i + \sum_{1 \leqslant i \leqslant k} \beta_i \cdot \tilde{\mathrm{u}}_i - \mathrm{w})/e \in \mathbb{F}_e^{n+m}$ and

$$x = \frac{\prod_{i \in I} x_i^{\nu_i} \prod_{1 \leqslant i \leqslant k} \tilde{x}_i^{\beta_i}}{g^{s'} \prod_{j=1}^{n} g_j^{w'_j} \prod_{j=1}^{m} h_j^{w'_{n+j}}} \bmod N. \tag{6}$$

Let $\mathrm{y} \in \mathbb{F}_e^n$ be the first n entries of w and $t = (s, x)$. The server sends $T = (T_1, T_2)$ as a proof of storage corresponding to the challenge set Q, where $T_1 = (\mathrm{y}, t)$ and $T_2 = \{(t_i, \Pi(i))\}_{i \in I}$.

Thus, during an audit, the server randomizes the linear combination of the challenged vectors with the help of a random linear combination of some predefined vectors prohibiting the TPA from gaining knowledge of the challenged vectors.

6 Enabling User-Anonymity for Shared Data in DSCS I

In a follow-up work of [16], Chen et al. [17] extend their SCS protocol for static data to a group setting where the users belonging to a group share the outsourced data (or file) and modify the same. In this scenario, the users in the group delegate the work of generating authentication tags for the vectors (of the file to be outsourced) to a security-mediator [50]. As the tags are generated by the mediator using the same secret key for all users, the cloud server cannot distinguish the users uploading (or updating) the file. The secret key is kept with the mediator only; thus, the users do not have the knowledge of the same. Moreover, a user obtains a tag corresponding to a vector w from the mediator in such a way that the mediator cannot gain knowledge of the content of w. A user outsources (or updates) the file using an anonymous channel [30]. Finally, every user in the group can perform an audit on the file kept in the cloud server. DSCS I can be extended to support anonymity for shared dynamic data as well. The procedures involved are described below.

– KeyGen$(1^\lambda, m, n)$: The mediator selects two random safe primes p, q of length $\lambda/2$ bits each and takes $N = pq$. It chooses another random prime e of length $\lambda + 1$ (in bits) and sets the file identifier $\mathtt{fid}$ to be equal to e. The mediator selects $g, g_1, \ldots, g_n, h_1, \ldots, h_m \xleftarrow{R} \mathbb{Z}_N^*$. The secret key sk is (p, q), and the public key pk consists of $(N, e, g, g_1, \ldots, g_n, h_1, \ldots, h_m, d_M, m, n)$. Initially, d_M is $\mathtt{null}$. Let $K = (sk, pk)$. We assume that every user in the group can update the public key efficiently.
– Blind(w): A user randomly selects vectors $\{\mathrm{w}_i\}_{1 \leqslant i \leqslant k}$ and $\{\beta_i\}_{1 \leqslant i \leqslant k}$ such that $\mathrm{w} = (\sum_{1 \leqslant i \leqslant k} \beta_i \cdot \mathrm{w}_i) \bmod e$. The user send the vectors $\{\mathrm{w}_i\}_{1 \leqslant i \leqslant k}$ to the mediator for their respective authentication tags. The mediator sends the tags $\{t_i\}_{1 \leqslant i \leqslant k}$ to the user.

– Unblind($\{\mathrm{w}_i, t_i\}_{1 \leqslant i \leqslant k}$): Due to the homomorphic property of the underlying signature scheme, the user computes the authentication tag t corresponding to the vector w in the same way as discussed in Section 3.3 (see the procedure Prove).
– Outsource($F, K,$ `fid`): Let the initial file F (associated with the identifier `fid`) consist of m vectors each of them having n blocks. We assume that each of these blocks is an element of $\mathbb{F}_e$. Then, for each $1 \leqslant i \leqslant m$, the i-th vector v_i is of the form $[v_{i1}, \ldots, v_{in}] \in \mathbb{F}_e^n$. For each vector v_i, a user obtains the corresponding authentication tag $t_i = (s_i, x_i)$ from the security-mediator (using the procedures Blind and Unblind described above) such that $s_i \xleftarrow{R} \mathbb{F}_e$ and

$$x_i^e = g^{s_i}(\prod_{j=1}^{n} g_j^{v_{ij}})h_i \bmod N. \tag{7}$$

The user constructs a rank-based authenticated skip list M on the authentication tags $\{t_i\}_{1 \leqslant i \leqslant m}$ and computes the metadata d_M (the label of the root node of M). Finally, the user updates d_M in the public key and uploads the file $F' = \{(\mathrm{v}_i, t_i)\}_{1 \leqslant i \leqslant m}$ along with M to the cloud server.
– InitUpdate($i,$ `updtype`$, d_M, pk,$ `fid`): The value of the variable `updtype` indicates whether the update is an insertion after or a modification of or the deletion of the i-th vector. The user performs one of the following operations depending on the value of `updtype`.
  1. If `updtype` is insertion, the user selects $h' \xleftarrow{R} \mathbb{Z}_N^*$ and generates the new vector-tag pair (v', t') using the procedures Blind and Unblind (with the help of the mediator). She runs ListInitUpdate on $(i,$ `updtype`$, d_M, t', m)$ and sends (h', v') to the server.
  2. If `updtype` is modification, the user generates the new vector-tag pair (v', t') using the procedures Blind and Unblind. Then, she runs ListInitUpdate($i,$ `updtype`$, d_M, t', m$) and sends v' to the server.
  3. If `updtype` is deletion, the user runs the procedure ListInitUpdate on $(i,$ `updtype`$, d_M, t', m)$, where t' is `null`.

  The user stores the value of the new metadata d'_M temporarily at her end.
– PerformUpdate($i,$ `updtype`$, F', M, h', \mathrm{v}', t', pk,$ `fid`): We assume that, for efficiency, the server keeps a local copy of the ordered list of h_j values for $1 \leqslant j \leqslant m$. Based on the value of `updtype`, the server performs one of the following operations.
  1. If `updtype` is insertion, the server sets $m = m + 1$, inserts h' in the $(i+1)$-th position in the list of h_j values (for $1 \leqslant j \leqslant m$) and inserts v' after the i-th vector. The server runs ListPerformUpdate on the input $(i,$ `updtype`$, t', M)$.
  2. If `updtype` is modification (h' is `null`), the server modifies the i-th vector to v' and runs the procedure ListPerformUpdate on $(i,$ `updtype`$, t', M)$.
  3. If `updtype` is deletion (h', v' and t' are `null`), the server sets $m = m - 1$, deletes the particular h_i value from the list of h_j values ($j \in [1, m]$) and runs ListPerformUpdate($i,$ `updtype`$,$ `null`$, M$).
– VerifyUpdate($i,$ `updtype`$, t', d'_M, \Pi, pk,$ `fid`): After receiving the proof from the server, the user performs ListVerifyUpdate($i,$ `updtype`$, t', d'_M, \Pi, m$). If the output of ListVerifyUpdate is 1, the user outputs 1 and updates the public key (the latest values of m, d_M and h_j for $j \in [1, m]$) accordingly. Otherwise, the user outputs 0.
– Challenge($pk, l,$ `fid`): A user performing an audit selects I, a random l-element subset of $[1, m]$. Then, she generates a challenge set $Q = \{(i, \nu_i)\}_{i \in I}$, where each $\nu_i \xleftarrow{R} \mathbb{F}_e$. The user sends the challenge set Q to the cloud server.
– Prove($Q, pk, F', M,$ `fid`): The cloud server, after receiving the challenge set $Q = \{(i, \nu_i)\}_{i \in I}$, computes $s = \sum_{i \in I} \nu_i s_i \bmod e$ and $s' = (\sum_{i \in I} \nu_i s_i - s)/e$. The server, for each $i \in I$, forms $\mathrm{u}_i = [\mathrm{v}_i\ \mathrm{e}_i] \in \mathbb{F}_e^{n+m}$ by augmenting the vector v_i with the unit coefficient vector e_i. Then, it computes $\mathrm{w} = \sum_{i \in I} \nu_i \cdot \mathrm{u}_i \bmod e \in \mathbb{F}_e^{n+m}$, $\mathrm{w}' = (\sum_{i \in I} \nu_i \cdot \mathrm{u}_i - \mathrm{w})/e \in \mathbb{F}_e^{n+m}$ and

$$x = \frac{\prod_{i \in I} x_i^{\nu_i}}{g^{s'} \prod_{j=1}^{n} g_j^{w'_j} \prod_{j=1}^{m} h_j^{w'_{n+j}}} \bmod N. \tag{8}$$

Let $\mathrm{y} \in \mathbb{F}_e^n$ be the first n entries of w and $t = (s, x)$. The server sends $T = (T_1, T_2)$ as a proof of storage corresponding to the challenge set Q, where $T_1 = (\mathrm{y}, t)$ and $T_2 = \{(t_i, \Pi(i))\}_{i \in I}$.

– Verify($Q, T, pk,$ `fid`): Using $Q = \{(i, \nu_i)\}_{i \in I}$ and $T = (\mathrm{y}, t)$ sent by the server, the user performing the audit constructs a vector $\mathrm{w} = [w_1, \ldots, w_n, w_{n+1}, \ldots, w_{n+m}] \in \mathbb{F}_e^{n+m}$, where the first n entries of w are the same as those of y and the $(n+i)$-th entry is ν_i if $i \in I$ (0 if $i \notin I$). Then, the user checks whether

$$x^e \stackrel{?}{=} g^s \prod_{j=1}^{n} g_j^{w_j} \prod_{j=1}^{m} h_j^{w_{n+j}} \bmod N. \tag{9}$$

She also verifies if, for each $i \in I$, $\Pi(i)$ is a valid proof (with respect to d_M) for t_i. The user outputs 1 if the proof passes all the verifications; she outputs 0, otherwise.

We have modified our basic DSCS I protocol (described in Section 3.3) to handle a group of users sharing their data without disclosing their identity to the cloud server. If we require a third party auditor (other than the users in the group) to audit on behalf of the group in a privacy-preserving fashion, it is not hard to see that the same changes (discussed in Section 5) can be applied in the group setting as well.

If we consider a procedure Update to be the triplet of the procedures (InitUpdate, PerformUpdate, VerifyUpdate), such an procedure Update must be performed atomically. Additionally, an update and an audit must not coincide in order to maintain the consistency of the outsourced data.

7 Performance Analysis of DSCS I

In this section, we discuss about the efficiency of our DSCS I protocol (described in Section 3.3) and compare this scheme with other existing SCS protocols achieving provable data possession guarantees. We also identify some limitations of an SNC-based SCS scheme (for static or dynamic data) compared to the DPDP I scheme (described in Appendix B.1).

7.1 Efficiency

The computational cost of the procedures in DSCS I is dominated by the cost of exponentiations (modulo N). To generate the value x in an authentication tag for each vector (in the procedure Outsource), the client has to perform a multi-exponentiation[3] and calculate the e-th root of the result (see Eqn. 1). The server requires two multi-exponentiations to calculate the value of x (see Eqn. 2 in the procedure Prove). To verify a proof using the procedure Verify, the verifier has to perform a multi-exponentiation and a single exponentiation (see Eqn. 3).

As mentioned in Section 3.2, due to the properties of a skip list [43], the size of each proof Π (related to the rank-based authenticated skip list), the time required to generate Π and the time required to verify Π are $O(\log m)$ with high probability.

7.2 Comparison among PDP Schemes

As DSCS I protocol provides provable data possession (PDP) guarantees, we compare our scheme with some other PDP schemes found in the literature. The comparison shown in Table 1 is done based on different parameters related to an audit. In Section 8, we propose a more efficient scheme suitable for append-only data that we mention as DSCS II in Table 1.

Now, we discuss about a few limitations of our DSCS I protocol compared to DPDP I (specifically), since both of them are secure in the *standard model*, handle *dynamic* data and offer *public verifiability*. In DSCS I, the audits are privacy-preserving, that is, a third party auditor (TPA) cannot gain knowledge of the data actually stored in the cloud server. Although the original DPDP I scheme does not offer privacy-preserving audits, this scheme can be modified to support the same (see Appendix B.2). The issues of our scheme compared to the modified DPDP I scheme are mentioned below.

[3] A naive way to compute a product of the form $\prod_{i=1}^{k} a_i^{e_i}$ in a finite group is to multiply the results of the individual exponentiations. There are better algorithms for computing such a multi-exponentiation [39].

Table 1. Comparison among Secure Cloud Storage Schemes Achieving PDP Guarantees

Secure cloud storage protocols	Type of data	Computation for verifier	Computation for server	Communication complexity	Publicly verifiable	Privacy-preserving audits	Security model
PDP [4]	Static	$O(1)$	$O(1)$	$O(1)$	Yes	No	RO†
Scalable PDP [5]	Dynamic‡	$O(1)$	$O(1)$	$O(1)$	No	No	RO
DPDP I [22]	Dynamic	$O(\log \tilde{m})$	$O(\log \tilde{m})$	$O(\log \tilde{m})$	Yes§	No	Standard
DPDP II [22]	Dynamic	$O(\log \tilde{m})$	$O(\tilde{m}^{\epsilon} \log \tilde{m})$⋆	$O(\log \tilde{m})$	Yes§	No	Standard
Wang et al. [52]	Dynamic	$O(\log \tilde{m})$	$O(\log \tilde{m})$	$O(\log \tilde{m})$	Yes	No	RO
Wang et al. [51]	Dynamic	$O(\log \tilde{m})$	$O(\log \tilde{m})$	$O(\log \tilde{m})$	Yes	Yes	RO
Chen et al. [16]	Static	$O(1)$	$O(1)$	$O(1)$	Yes	Yes	Standard
DSCS I (in this work)	Dynamic	$O(\log m)$	$O(\log m)$	$O(\log m)$	Yes	Yes	Standard
Modified DPDP I⋆⋆ [46]	Dynamic	$O(\log \tilde{m})$	$O(\log \tilde{m})$	$O(\log \tilde{m})$	Yes	Yes	Standard
DSCS II (in this work)	Dynamic¶	$O(\log m)$	$O(\log m)$	$O(\log m)$	Yes	Yes	RO

For simplicity, we exclude the security parameter λ from complexity parameters (for an audit). The value $\tilde{m}$ denotes the number of segments the data file is divided in (such that an authentication tag is associated with each segment). For example, $\tilde{m} = m$ in our DSCS (I and II) schemes, where m denotes the number of vectors. The term $O(\tilde{n})$ is added implicitly to each complexity parameter, where $\tilde{n}$ is the size of each segment. For example, $\tilde{n} = n$ in DSCS I and II, where a vector having n blocks is considered to be a segment. For all the schemes, the storage at the verifier side is $O(1)$, and the storage at the server side is $O(|F'|)$ where F' is the outsourced file. If l is the cardinality of the challenge set and the server corrupts β fraction of the file, the detection probability $p_{detect} = 1 - (1-\beta)^l$ for all the schemes (except, in DPDP II, $p_{detect} = 1 - (1-\beta)^{\Omega(\log \tilde{m})}$).

† RO denotes the random oracle model [7].

‡ Scalable PDP scheme supports deletion, modification and append only for a predefined number of times; insertion is not supported in this scheme.

§ A small change (making the latest values of $d_{\tilde{M}}$ and $\tilde{m}$ public) is required in the original scheme (see Section B.1).

⋆ ϵ is a constant such that $0 < \epsilon < 1$.

⋆⋆ In the preliminary version of this paper [46], we modify DPDP I [22] to make its audits privacy-preserving. The modified DPDP I scheme is described in Appendix B.2.

¶ DSCS II supports only append and modification; arbitrary insertion (or deletion) is not supported.

1. The size of the public key is $O(m+n)$ in DSCS I. On the other hand, the size of the public key in the modified DPDP I scheme is constant.
2. The authentication tags in DSCS I are of the form (s, x), where $s \in \mathbb{F}_e$ and $x \in \mathbb{Z}_N^*$. An authentication tag in the modified DPDP I scheme is an element of $\mathbb{Z}_N^*$. Thus, the size of a tag in DSCS I is larger than that in the modified DPDP I scheme by $\lambda + 1$ bits (as e is a $(\lambda+1)$-bit prime).
3. In DSCS I, the value of (d_M, m) and the h_i values in the public key must be changed for each insertion or deletion (only change in d_M is required for modification), whereas only the value of $(d_{\tilde{M}}, \tilde{m})$ needs to be changed in the modified DPDP I scheme. However, if the server keeps a local copy of the public key (an ordered list containing h_i values for $i \in [1, m]$), then small changes are required at the server side. The server inserts the new h value (sent by the client) in $(i+1)$-th position in the list (for insertion) or discards the i-th h value (for deletion).

Thus, the proposed DSCS I scheme suffers from the limitations mentioned above. We note that the existing SCS protocol for static data [16] based on the same SNC protocol [12] also suffers from the first two of these limitations. However, in this work, we explore if a secure cloud storage protocol for dynamic data can be constructed from a secure network coding protocol. A more efficient (in terms of the size of the public key or the size of an

authentication tag) SNC protocol can lead us to the construction of a more efficient DSCS protocol in future. In the following section, we propose another secure cloud storage protocol (DSCS II) for append-only data that is much more efficient than DSCS I.

8 More Efficient Solutions for Append-only Data

Although generic dynamic data are useful, append-only data find numerous applications as well. These primarily include archival data from different sources where data are appended to the existing datasets. For example, data obtained from closed circuit television (CCTV) camera, monetary transactions in banks, medical history of patients — all must be kept intact with append being the only possible update. On the other hand, various cloud service providers like Amazon Web Services (AWS) use Hadoop Distributed File System (HDFS [24]) to store large volume of data. In HDFS, data blocks are added by an application to a new file. After closing the file, the blocks cannot be removed or modified further; blocks can only be appended by reopening the file. Append-only data are also useful for maintaining log structures (e.g., certificates are stored using append-only log structures in certificate transparency schemes [34,45]).

In this section, we construct a more efficient DSCS scheme (DSCS II) for append-only data using the SNC protocol proposed by Boneh et al. [8] that is not suitable for constructing a DSCS scheme for generic dynamic data (see Section 3.1). First, we give a brief overview of the SNC protocol and discuss a key property of it that allows us to construct DSCS II based on the protocol. Then, we provide the construction of the DSCS II scheme.

8.1 Homomorphic Network Coding Signature Scheme Proposed by Boneh et al.

Boneh et al. [8] propose a homomorphic network coding signature scheme secure in the random oracle model [7] under the co-computational Diffie Hellman (co-CDH) assumption. We briefly describe the procedures involved in this scheme. The notations are the same as those discussed in Section 2.2.

- KeyGen$(1^\lambda, m, n)$: Let $\mathcal{G} = (G_1, G_2, G_T, e, \psi)$ be a bilinear group tuple, where G_1, G_2 and G_T are multiplicative cyclic groups of prime order $p > 2^\lambda$, and the functions $e : G_1 \times G_2 \to G_T$ (bilinear map) and $\psi : G_2 \to G_1$ are efficiently computable. Choose $g_1, \ldots, g_n \xleftarrow{R} G_1 \backslash \{1\}$, $h \xleftarrow{R} G_2 \backslash \{1\}$ and $\alpha \xleftarrow{R} \mathbb{F}_p$. Take $u = h^\alpha$. Let $H : \mathbb{Z} \times \mathbb{Z} \to G_1$ be a hash function considered to be a random oracle. The public key is $pk = (\mathcal{G}, H, g_1, \ldots, g_n, h, u)$, and the private key is $sk = \alpha$.
- TagGen$(V, sk, m, n, \texttt{fid})$: Given the secret key sk, a linear subspace $V \subset \mathbb{F}_p^{n+m}$ spanned by the augmented vectors $\mathrm{u}_1, \mathrm{u}_2, \ldots, \mathrm{u}_m$ and a random file identifier $\texttt{fid} \in \{0,1\}^\lambda$, the sender outputs the signature $t_i = \left(\prod_{j=1}^{n} g_j^{u_{ij}} \prod_{j=1}^{m} H(\texttt{fid}, j)^{u_{i(n+j)}} \right)^\alpha$ for the i-th basis vector $\mathrm{u}_i = [u_{i1}, u_{i2}, \ldots, u_{i(n+m)}] \in \mathbb{F}_p^{n+m}$ for each $i \in [1, m]$.
- Combine$(\{\mathrm{y}_i, t_i, \nu_i\}_{1 \leqslant i \leqslant l}, pk, m, n, \texttt{fid})$: Given the public key pk, the file identifier $\texttt{fid}$ and l tuples (each consisting of a vector $\mathrm{y}_i \in \mathbb{F}_p^{n+m}$, a coefficient $\nu_i \in \mathbb{F}_p$ and a signature t_i), an intermediate node outputs the signature $t = \prod_{i=1}^{l} t_i^{\nu_i}$ for another vector $\mathrm{w} = \sum_{i=1}^{l} \nu_i \cdot \mathrm{y}_i \in \mathbb{F}_p^{n+m}$.
- Verify$(\mathrm{w}, t, pk, m, n, \texttt{fid})$: Given the public key pk, the unique file identifier $\texttt{fid}$, a signature t and a vector $\mathrm{w} = [w_1, w_2, \ldots, w_{n+m}] \in \mathbb{F}_p^{n+m}$, an intermediate node or the receiver node checks whether

$$e(t, h) \stackrel{?}{=} e\left(\prod_{j=1}^{n} g_j^{w_j} \prod_{j=1}^{m} H(\texttt{fid}, j)^{w_{n+j}}, u \right).$$

  If the equality holds, it outputs 1; it outputs 0, otherwise.

We recall that in a secure cloud storage protocol (using secure network coding), the client divides the file F associated with $\texttt{fid}$ into m vectors each of them having n blocks. The i-th vector v_i is of the form $[v_{i1}, \ldots, v_{in}] \in$

$\mathbb{F}^n, \forall i \in [1, m]$. For each vector v_i, the client forms $\mathrm{u}_i = [\mathrm{v}_i\ \mathrm{e}_i] \in \mathbb{F}^{n+m}$ by augmenting the vector v_i with the unit coefficient vector e_i. If we use the current SNC protocol [8], the client runs TagGen$(V, sk, m, n, \mathtt{fid})$ to produce, for each $i \in [1, m]$, a signature (authentication tag)

$$t_i = \left(\prod_{j=1}^{n} g_j^{u_{ij}} \prod_{j=1}^{m} H(\mathtt{fid}, j)^{u_{i(n+j)}} \right)^{\alpha}$$
$$= \left(H(\mathtt{fid}, i) \prod_{j=1}^{n} g_j^{u_{ij}} \right)^{\alpha} = \left(H(\mathtt{fid}, i) \prod_{j=1}^{n} g_j^{v_{ij}} \right)^{\alpha}, \tag{10}$$

for the vector u_i. We observe that the vector index i is embedded in the tag corresponding to the i-th vector. Therefore, the scheme is not suitable for construction of a secure cloud storage for dynamic (in generic sense) data as mentioned in Section 3.1.

8.2 DSCS II: An Efficient Solution for Append-only Data

From the previous section, we note that the authentication tags in the SNC protocol [8] are independent of the value m (see Eqn. 10). The size of the public key also does not depend on m. This makes the SNC protocol suitable for constructing a more efficient DSCS scheme. However, as each authentication tag embeds the index of the respective vector, we cannot insert or delete at arbitrary positions of the data file. Since the value of the index i of the vector to be inserted only increases for append-only data, the SNC protocol provides an efficient DSCS protocol (DSCS II). We observe that, for the same reason, the SNC protocol proposed by Gennaro et al. [26] can also be used for such a construction. Although we construct DSCS II to handle append-only data, it supports modifications of existing vectors trivially. DSCS II consists of the following procedures.

- KeyGen$(1^\lambda, m, n)$: Let $\mathcal{G} = (G_1, G_2, G_T, e, \psi)$ be a bilinear group tuple, where G_1, G_2 and G_T are multiplicative cyclic groups of prime order $p > 2^\lambda$, and the functions $e : G_1 \times G_2 \to G_T$ (bilinear map) and $\psi : G_2 \to G_1$ are efficiently computable. The client selects $g_1, \ldots, g_n \xleftarrow{R} G_1 \backslash \{1\}$, $h \xleftarrow{R} G_2 \backslash \{1\}$ and $\alpha \xleftarrow{R} \mathbb{F}_p$. She takes $u = h^\alpha$ and chooses a random file identifier $\mathtt{fid} \in \{0,1\}^\lambda$. Let $H : \mathbb{Z} \times \mathbb{Z} \to G_1$ be a hash function considered to be a random oracle. The public key is $pk = (\mathcal{G}, H, g_1, \ldots, g_n, h, u, d_M, m, n)$, and the private key is $sk = \alpha$. Initially, d_M is $\mathtt{null}$. Let $K = (sk, pk)$.
- Outsource$(F, K, \mathtt{fid})$: The file F (associated with the identifier $\mathtt{fid}$) consists of m vectors each of them having n blocks. We assume that each of these blocks is an element of $\mathbb{F}_p$. Then, for each $1 \leqslant i \leqslant m$, the i-th vector v_i is of the form $[v_{i1}, \ldots, v_{in}] \in \mathbb{F}_p^n$. For each vector v_i, the client computes the authentication tag

$$t_i = \left(H(\mathtt{fid}, i) \prod_{j=1}^{n} g_j^{v_{ij}} \right)^{\alpha} \tag{11}$$

  as shown in Eqn. 10. The client constructs a rank-based authenticated skip list M on the authentication tags $\{t_i\}_{1 \leqslant i \leqslant m}$ and computes the metadata d_M (the label of the root node of M). Finally, the client updates d_M in the public key and uploads the file $F' = \{(\mathrm{v}_i, t_i)\}_{1 \leqslant i \leqslant m}$ along with M to the cloud server.
- InitUpdate$(i, \mathtt{updtype}, pk, \mathtt{fid})$: The value of the variable $\mathtt{updtype}$ indicates whether the update is a modification of the i-th vector or an append at the end of the file. The client performs one of the following operations depending on the value of $\mathtt{updtype}$.
  1. If $\mathtt{updtype}$ is append, the client generates the new vector-tag pair (v', t') and sets $\mathtt{UPDTYPE}$ to be insertion. She runs ListInitUpdate on $(m, \mathtt{UPDTYPE}, d_M, t', m)$ and sends v' to the server.
  2. If $\mathtt{updtype}$ is modification, the client generates the new vector-tag pair (v', t'). Then, the client runs ListInitUpdate$(i, \mathtt{updtype}, d_M, t', m)$ and sends v' to the server.

  The client stores the value of the new metadata d'_M temporarily at her end.
- PerformUpdate$(i, \mathtt{updtype}, F', M, h', \mathrm{v}', t', pk, \mathtt{fid})$: Based on the value of $\mathtt{updtype}$, the server performs one of the following operations.

1. If updtype is append, the server inserts v′ after the m-th vector and sets UPDTYPE to be insertion. The server runs ListPerformUpdate on the input $(m, \texttt{UPDTYPE}, t', M)$ and sets $m = m + 1$.
2. If updtype is modification, the server modifies the i-th vector to v′ and runs the procedure ListPerformUpdate on $(i, \texttt{updtype}, t', M)$.

– VerifyUpdate$(i, \texttt{updtype}, t', d'_M, \Pi, pk, \texttt{fid})$: After receiving the proof from the server, the client performs one of the following operations based on the value of updtype.
1. If updtype is append, the client sets UPDTYPE to be insertion and executes the procedure ListVerifyUpdate on $(i, \texttt{UPDTYPE}, t', d'_M, \Pi, m)$.
2. If updtype is modification, the client executes ListVerifyUpdate$(i, \texttt{updtype}, t', d'_M, \Pi, m)$.

If the output of ListVerifyUpdate is 1, the client outputs 1 and updates her public key (the latest values of m and d_M) accordingly. Otherwise, the client outputs 0.

– Challenge$(pk, l, \texttt{fid})$: During an audit, the verifier selects I, a random l-element subset of $[1, m]$. Then, she generates a challenge set $Q = \{(i, \nu_i)\}_{i \in I}$, where each $\nu_i \xleftarrow{R} \mathbb{F}_p$. The verifier sends the challenge set Q to the cloud server.

– Prove$(Q, pk, F', M, \texttt{fid})$: Upon receiving the challenge set $Q = \{(i, \nu_i)\}_{i \in I}$, the cloud server, for each $i \in I$, forms $\mathrm{u}_i = [\mathrm{v}_i \ \mathrm{e}_i] \in \mathbb{F}_p^{n+m}$ by augmenting the vector v_i with the unit coefficient vector e_i. Then, it computes the authentication tag

$$t = \prod_{i=1}^{l} t_i^{\nu_i} \tag{12}$$

for another vector $\mathrm{w} = \sum_{i=1}^{l} \nu_i \cdot \mathrm{u}_i \in \mathbb{F}_p^{n+m}$. Let $\mathrm{y} \in \mathbb{F}_p^n$ be the first n entries of w. The server sends $T = (T_1, T_2)$ as a proof of storage corresponding to the challenge set Q, where $T_1 = (\mathrm{y}, t)$ and $T_2 = \{(t_i, \Pi(i))\}_{i \in I}$.

– Verify$(Q, T, pk, \texttt{fid})$: Using $Q = \{(i, \nu_i)\}_{i \in I}$ and $T = (\mathrm{y}, t)$ sent by the server, the verifier constructs a vector $\mathrm{w} = [w_1, \ldots, w_n, w_{n+1}, \ldots, w_{n+m}] \in \mathbb{F}_p^{n+m}$, where the first n entries of w are the same as those of y and the $(n+i)$-th entry is ν_i if $i \in I$ (0 if $i \notin I$). Then, the verifier checks whether

$$e(t, h) \stackrel{?}{=} e\left(\prod_{j=1}^{n} g_j^{w_j} \prod_{j=1}^{m} H(\texttt{fid}, j)^{w_{n+j}}, u\right). \tag{13}$$

She also verifies if, for each $i \in I$, $\Pi(i)$ is a valid proof (with respect to d_M) for t_i. The verifier outputs 1 if the proof passes all the verifications; she outputs 0, otherwise.

The secure cloud storage scheme (DSCS II) described above supports only append at the end of the data file and modification of any existing vector in the file. The scheme is publicly verifiable in that anyone with the knowledge of the public key can perform an audit. The scheme is secure in the random oracle model (according to Definition 3 in Section 4). The security proof of DSCS II is the same as that of our DSCS I protocol (see Theorem 1) except that the guarantee of authenticity comes from the security of the underlying SNC protocol [8] that is secure in the random oracle model. Moreover, the audits in this scheme can be made privacy-preserving in the same way as discussed in Section 5.

Efficiency In the procedure Outsource, the client has to perform a multi-exponentiation to generate the value of an authentication tag for each vector t (see Eqn. 11). The server requires one multi-exponentiation to calculate the value of t (see Eqn. 12 in the procedure Prove). The verifier has to perform two multi-exponentiations and two pairing operations (see Eqn. 13) to verify a proof using the procedure Verify. Due to the properties of a skip list [43], the size of each proof Π (related to the rank-based authenticated skip list), the time required to generate Π and the time required to verify Π are $O(\log m)$ with high probability. Different parameters of the scheme (DSCS II) related to an audit are shown in Table 1 (Section 7.2).

An authentication tag t in DSCS II belongs to G_1, and thus, is of size 2λ bits for Type 2 pairings [25]. The DSCS II scheme for append-only data overcomes some of the limitations of our DSCS I protocol described in Section 3.3 as follows.

1. In DSCS II, the size of the public key is $O(n)$ (which is $O(m+n)$ in DSCS I), where n ($\ll m$) is fixed during the setup and kept unchanged during the execution of the protocol.
2. In DSCS II, only the value of (d_M, m) needs to be changed for an append operation that is similar to DPDP I [22]. On the other hand, the h_i values in the public key, in addition to the value of (d_M, m), need to be changed for each insertion or deletion in DSCS I.

9 Conclusion

In this work, we have proposed a DSCS protocol (DSCS I) based on an SNC protocol. To the best of our knowledge, this is the first SNC-based DSCS protocol that is secure in the standard model and enjoys public verifiability. We have also discussed about some challenges while constructing, in general, an efficient DSCS protocol from an SNC protocol. We have shown that DSCS I can be extended to support privacy-preserving audits and to preserve anonymity for shared data where the cloud server cannot distinguish the users in a group updating the shared data. We have analyzed the efficiency of our DSCS construction and compare it with other existing secure cloud storage protocols achieving the guarantees of provable data possession. We have also identified some limitations of an SNC-based secure cloud storage protocol. However, some of these limitations follow from the underlying SNC protocols used. A more efficient SNC protocol can give us a DSCS protocol with a better efficiency. Finally, we have constructed another DSCS II scheme for append-only data that overcomes some of the limitations of DSCS I.

References

1. Agrawal, S., Boneh, D.: Homomorphic MACs: MAC-based integrity for network coding. In: Applied Cryptography and Network Security - ACNS 2009. (2009) 292–305
2. Ahlswede, R., Cai, N., Li, S.R., Yeung, R.W.: Network information flow. IEEE Transactions on Information Theory **46**(4) (2000) 1204–1216
3. Armknecht, F., Bohli, J., Karame, G.O., Liu, Z., Reuter, C.A.: Outsourced proofs of retrievability. In: ACM Conference on Computer and Communications Security, CCS 2014. (2014) 831–843
4. Ateniese, G., Burns, R.C., Curtmola, R., Herring, J., Kissner, L., Peterson, Z.N.J., Song, D.X.: Provable data possession at untrusted stores. In: ACM Conference on Computer and Communications Security, CCS 2007. (2007) 598–609
5. Ateniese, G., Pietro, R.D., Mancini, L.V., Tsudik, G.: Scalable and efficient provable data possession. In: International Conference on Security and Privacy in Communication Networks, SECURECOMM 2008. (2008) 9
6. Attrapadung, N., Libert, B.: Homomorphic network coding signatures in the standard model. In: Public Key Cryptography - PKC 2011. (2011) 17–34
7. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: ACM Conference on Computer and Communications Security, CCS 1993. (1993) 62–73
8. Boneh, D., Freeman, D.M., Katz, J., Waters, B.: Signing a linear subspace: Signature schemes for network coding. In: Public Key Cryptography - PKC 2009. (2009) 68–87
9. Bowers, K.D., Juels, A., Oprea, A.: HAIL: A high-availability and integrity layer for cloud storage. In: ACM Conference on Computer and Communications Security, CCS 2009. (2009) 187–198
10. Bowers, K.D., Juels, A., Oprea, A.: Proofs of retrievability: Theory and implementation. In: ACM Cloud Computing Security Workshop, CCSW 2009. (2009) 43–54
11. Cash, D., Küpçü, A., Wichs, D.: Dynamic proofs of retrievability via oblivious RAM. In: Advances in Cryptology - EUROCRYPT 2013, Springer Berlin Heidelberg (2013) 279–295
12. Catalano, D., Fiore, D., Warinschi, B.: Efficient network coding signatures in the standard model. In: Public Key Cryptography - PKC 2012. (2012) 680–696
13. Chandran, N., Kanukurthi, B., Ostrovsky, R.: Locally updatable and locally decodable codes. In: Theory of Cryptography Conference, TCC 2014. (2014) 489–514
14. Charles, D.X., Jain, K., Lauter, K.E.: Signatures for network coding. International Journal of Information and Coding Theory **1**(1) (2009) 3–14
15. Chen, B., Curtmola, R., Ateniese, G., Burns, R.C.: Remote data checking for network coding-based distributed storage systems. In: ACM Cloud Computing Security Workshop, CCSW 2010. (2010) 31–42
16. Chen, F., Xiang, T., Yang, Y., Chow, S.S.M.: Secure cloud storage meets with secure network coding. In: IEEE Conference on Computer Communications, INFOCOM 2014. (2014) 673–681

17. Chen, F., Xiang, T., Yang, Y., Chow, S.S.M.: Secure cloud storage meets with secure network coding. IEEE Transactions on Computers **65**(6) (2016) 1936–1948
18. Curtmola, R., Khan, O., Burns, R.C., Ateniese, G.: MR-PDP: Multiple-replica provable data possession. In: IEEE International Conference on Distributed Computing Systems, ICDCS 2008. (2008) 411–420
19. Dimakis, A.G., Godfrey, B., Wu, Y., Wainwright, M.J., Ramchandran, K.: Network coding for distributed storage systems. IEEE Transactions on Information Theory **56**(9) (2010) 4539–4551
20. Dodis, Y., Vadhan, S.P., Wichs, D.: Proofs of retrievability via hardness amplification. In: Theory of Cryptography Conference, TCC 2009. (2009) 109–127
21. Erway, C.C., Küpçü, A., Papamanthou, C., Tamassia, R.: Dynamic provable data possession. In: ACM Conference on Computer and Communications Security, CCS 2009. (2009) 213–222
22. Erway, C.C., Küpçü, A., Papamanthou, C., Tamassia, R.: Dynamic provable data possession. ACM Transactions on Information and System Security **17**(4) (2015) 15
23. Esiner, E., Küpçü, A., Özkasap, Ö.: Analysis and optimization on flexdpdp: A practical solution for dynamic provable data possession. In: Intelligent Cloud Computing - First International Conference, ICC 2014. (2014) 65–83
24. Foundation, T.A.S.: Apache hadoop (June 2016)
25. Galbraith, S.D., Paterson, K.G., Smart, N.P.: Pairings for cryptographers. Discrete Applied Mathematics **156**(16) (September 2008) 3113–3121
26. Gennaro, R., Katz, J., Krawczyk, H., Rabin, T.: Secure network coding over the integers. In: Public Key Cryptography - PKC 2010. (2010) 142–160
27. Goodrich, M.T., Tamassia, R., Schwerin, A.: Implementation of an authenticated dictionary with skip lists and commutative hashing. In: DARPA Information Survivability Conference and Exposition (DISCEX) II. (2001) 68–82
28. Ho, T., Koetter, R., Médard, M., Karger, D.R., Effros, M.: The benefits of coding over routing in a randomized setting. In: IEEE International Symposium on Information Theory - ISIT 2003. (2003) 442
29. Ho, T., Médard, M., Koetter, R., Karger, D.R., Effros, M., Shi, J., Leong, B.: A random linear network coding approach to multicast. IEEE Transactions on Information Theory **52**(10) (2006) 4413–4430
30. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Sahai, A.: Cryptography from anonymity. In: IEEE Symposium on Foundations of Computer Science - FOCS 2006. (2006) 239–248
31. Juels, A., Kaliski, Jr., B.S.: PORs: Proofs of retrievability for large files. In: ACM Conference on Computer and Communications Security, CCS 2007. (2007) 584–597
32. Krohn, M.N., Freedman, M.J., Mazières, D.: On-the-fly verification of rateless erasure codes for efficient content distribution. In: IEEE Symposium on Security and Privacy - S&P 2004. (2004) 226–240
33. Küpçü, A.: Official arbitration with secure cloud storage application. The Computer Journal **58**(4) (2015) 831–852
34. Laurie, B., Langley, A., Kasper, E.: Certificate transparency (June 2013)
35. Le, A., Markopoulou, A.: NC-audit: Auditing for network coding storage. In: International Symposium on Network Coding, NetCod 2012. (2012) 155–160
36. Li, S.R., Yeung, R.W., Cai, N.: Linear network coding. IEEE Transactions on Information Theory **49**(2) (2003) 371–381
37. MacWilliams, F.J., Sloane, N.J.A.: The Theory of Error-Correcting Codes. North-Holland Publishing Company (1977)
38. Merkle, R.C.: A digital signature based on a conventional encryption function. In: Advances in Cryptology - CRYPTO 1987. (1987) 369–378
39. Möller, B.: Algorithms for multi-exponentiation. In: Selected Areas in Cryptography - SAC 2001. (2001) 165–180
40. Naor, M., Rothblum, G.N.: The complexity of online memory checking. Journal of the ACM **56**(1) (February 2009) 2:1–2:46
41. Omote, K., Phuong, T.T.: DD-POR: Dynamic operations and direct repair in network coding-based proof of retrievability. In: Computing and Combinatorics - 21st International Conference, COCOON 2015. (2015) 713–730
42. Papamanthou, C., Tamassia, R., Triandopoulos, N.: Authenticated hash tables. In: ACM Conference on Computer and Communications Security, CCS 2008. (2008) 437–448
43. Pugh, W.: Skip lists: A probabilistic alternative to balanced trees. Communications of the ACM **33**(6) (1990) 668–676
44. Reed, I.S., Solomon, G.: Polynomial codes over certain finite fields. Journal of the Society for Industrial and Applied Mathematics **8**(2) (1960) 300–304
45. Ryan, M.D.: Enhanced certificate transparency and end-to-end encrypted mail. In: 21st Annual Network and Distributed System Security Symposium - NDSS 2014. (2014)
46. Sengupta, B., Ruj, S.: Publicly verifiable secure cloud storage for dynamic data using secure network coding. In: ACM Asia Conference on Computer and Communications Security, AsiaCCS 2016. (2016) 107–118
47. Shacham, H., Waters, B.: Compact proofs of retrievability. Journal of Cryptology **26**(3) (2013) 442–483
48. Shi, E., Stefanov, E., Papamanthou, C.: Practical dynamic proofs of retrievability. In: ACM Conference on Computer and Communications Security, CCS 2013. (2013) 325–336
49. Stefanov, E., van Dijk, M., Juels, A., Oprea, A.: Iris: A scalable cloud file system with efficient integrity checks. In: Annual Computer Security Applications Conference, ACSAC 2012. (2012) 229–238

50. Wang, B., Chow, S.S.M., Li, M., Li, H.: Storing shared data on the cloud via security-mediator. In: International Conference on Distributed Computing Systems, ICDCS 2013. (2013) 124–133
51. Wang, C., Chow, S.S.M., Wang, Q., Ren, K., Lou, W.: Privacy-preserving public auditing for secure cloud storage. IEEE Transactions on Computers **62**(2) (2013) 362–375
52. Wang, Q., Wang, C., Ren, K., Lou, W., Li, J.: Enabling public auditability and data dynamics for storage security in cloud computing. IEEE Transactions on Parallel and Distributed Systems **22**(5) (2011) 847–859

A General Construction of an SCS Protocol for Static Data from an SNC Protocol

Chen et al. [16] propose a generic construction of a secure cloud storage protocol for static data from a secure network coding protocol. They consider the data file F to be stored in the server to be a collection of m vectors (or packets) each of which consists of n blocks. The underlying idea is to store these vectors (without augmenting them with unit vectors) along with their authentication tags in the server. During an audit, the client sends an l-element subset of the set of indices $\{1, 2, \ldots, m\}$ to the server. The server augments those vectors with the corresponding unit vectors, combines them linearly in an authenticated fashion and sends the output vector along with its tag to the client. Finally, the client verifies the authenticity of the received tag against the received vector. Thus, the server acts as an intermediate node, and the client acts as both the sender and the receiver (or the next intermediate router). We briefly discuss the procedures involved in the general construction below.

- SSCS.KeyGen$(1^\lambda, m, n)$: Initially, the client executes SNC.KeyGen$(1^\lambda, m, n)$ to generate a secret key-public key pair $K = (sk, pk)$.
- SSCS.Outsource$(F, K, m, n, \texttt{fid})$: The file F associated with a random file identifier $\texttt{fid}$ consists of m vectors each of them having n blocks. We assume that each of these blocks is an element of $\mathbb{F}$. Then, for each $1 \leqslant i \leqslant m$, the i-th vector v_i is of the form $[v_{i1}, \ldots, v_{in}] \in \mathbb{F}^n$. For each vector v_i, the client forms $\mathrm{u}_i = [\mathrm{v}_i\ \mathrm{e}_i] \in \mathbb{F}^{n+m}$ by augmenting the vector v_i with the unit coefficient vector e_i. Let $V \subset \mathbb{F}^{n+m}$ be the linear subspace spanned by $\mathrm{u}_1, \mathrm{u}_2, \ldots, \mathrm{u}_m$. The client runs SNC.TagGen$(V, sk, m, n, \texttt{fid})$ to produce an authentication tag t_i for the i-th vector u_i for each $1 \leqslant i \leqslant m$. Finally, the client uploads the file $F' = \{(\mathrm{v}_i, t_i)\}_{1 \leqslant i \leqslant m}$ to the server.
- SSCS.Challenge$(pk, l, m, n, \texttt{fid})$: During an audit, the verifier selects I, a random l-element subset of $[1, m]$. Then, she generates a challenge set $Q = \{(i, \nu_i)\}_{i \in I}$, where each $\nu_i \xleftarrow{R} \mathbb{F}$. The verifier sends the challenge set Q to the server.
- SSCS.Prove$(Q, pk, F', m, n, \texttt{fid})$: Upon receiving the challenge set $Q = \{(i, \nu_i)\}_{i \in I}$ for the file identifier $\texttt{fid}$, the cloud server, for each $i \in I$, forms $\mathrm{u}_i = [\mathrm{v}_i\ \mathrm{e}_i] \in \mathbb{F}^{n+m}$ by augmenting the vector v_i with the unit coefficient vector e_i. Then, the cloud server runs SNC.Combine$(\{\mathrm{u}_i, t_i, \nu_i\}_{i \in I}, pk, m, n, \texttt{fid})$ to produce another vector $\mathrm{w} \in \mathbb{F}^{n+m}$ (along with its authentication tag t) such that $\mathrm{w} = \sum_{i=1}^{l} \nu_i \cdot \mathrm{u}_i$. Let $\mathrm{y} \in \mathbb{F}^n$ be the first n entries of w. The server sends $T = (\mathrm{y}, t)$ to the verifier as a proof of storage corresponding to the challenge set Q.
- SSCS.Verify$(Q, T, K, m, n, \texttt{fid})$: The verifier uses $Q = \{(i, \nu_i)\}_{i \in I}$ and $T = (\mathrm{y}, t)$ to reconstruct the vector $\mathrm{w} \in \mathbb{F}^{n+m}$, where the first n entries of w are the same as those of y and the $(n+i)$-th entry is ν_i if $i \in I$ (0 if $i \notin I$). The verifier runs SNC.Verify$(\mathrm{w}, t, K, m, n, \texttt{fid})$ and returns the output of the procedure SNC.Verify.

B DPDP I: A Dynamic Provable Data Possession Scheme

Erway et al. [22,21] propose two efficient and fully dynamic provable data possession schemes: DPDP I (based on rank-based authenticated skip lists) and DPDP II (based on rank-based RSA trees). We consider only the DPDP I scheme here.

B.1 Overview of DPDP I

Let there be a key generation procedure KeyGen that produces a public key $pk = (N, g)$, where $N = pq$ is a product of two large primes and g is an element of $\mathbb{Z}_N^*$ with large order. Suppose the initial data file consists of $\tilde{m}$

segments $b_1, b_2, \ldots, b_{\tilde{m}}$. For each segment b, the client computes a tag $\mathcal{T}(b) = g^b \bmod N$. Now, the client builds a rank-based authenticated skip list $\tilde{M}$ on the tags of the segments and uploads the data, tags and the skip list to the cloud server. The insertion, deletion and modification operations are performed in a similar fashion as discussed in Section 5. There is no secret key involved in the DPDP I scheme. Although Erway et al. do not claim explicitly the public verifiability of the DPDP I scheme, we observe that the scheme can be made publicly verifiable by simply making the metadata $d_{\tilde{M}}$ of the up-to-date skip list and the value $\tilde{m}$ public (see the footnote in Section 2.3).

During an audit, the verifier selects I, a random l-element subset of $\{1, 2, \ldots, \tilde{m}\}$, and generates a challenge set $Q = \{(i, \nu_i)\}_{i \in I}$, where each ν_i is a random value. The verifier sends the challenge set Q to the server. The server computes an aggregated segment $B = \sum_{i \in I} \nu_i b_i$ and sends $\{\mathcal{T}(b_i)\}_{i \in I}$, B and proofs $\{\Pi(i)\}_{i \in I}$ (see Section 3.2) to the verifier. The verifier computes $\mathcal{T} = \prod_{i \in I} \mathcal{T}(b_i)^{\nu_i}$. Finally, the verifier accepts the proof if and only if the following two conditions hold: $\Pi(i)$ is a valid proof for each $i \in I$ and $\mathcal{T} = g^B \bmod N$.

B.2 Modified DPDP I to Support Privacy-Preserving Audits

The secure cloud storage scheme for dynamic data discussed in Section 5 offers privacy-preserving audits where a third party auditor (TPA) cannot learn about the actual data while auditing. Let us investigate whether the DPDP I scheme provides this facility.

As in the original scheme (see Appendix B.1), the server sends the aggregated segment $B = \sum_{i \in I} \nu_i b_i$ to the verifier (or TPA) where $|I| = l$. Now, a TPA can obtain the b_i values by solving a system of linear equations. Therefore, the audits in the original scheme are not privacy-preserving. However, it is not hard to make these audits privacy-preserving. We modify the procedures involved in an audit as follows. As before, the verifier sends the challenge set $Q = \{(i, \nu_i)\}_{i \in I}$ to the server. The server computes an aggregated segment $B = \sum_{i \in I} \nu_i b_i$. Now, the server chooses a random value r, and it computes $B' = B + r$ and $R = g^r \bmod N$. The server sends $\{\mathcal{T}(b_i)\}_{i \in I}$, B', R and proofs $\{\Pi(i)\}_{i \in I}$ to the verifier. The verifier computes $\mathcal{T} = R \prod_{i \in I} \mathcal{T}(b_i)^{\nu_i}$. Finally, the verifier accepts the proof if and only if the following two conditions hold: $\Pi(i)$ is a valid proof for each $i \in I$ and $\mathcal{T} = g^{B'} \bmod N$.

As discussed in Appendix B.1, in order to make the scheme publicly verifiable, the client includes the pair $(d_{\tilde{M}}, \tilde{m})$ in her public key and updates it after every authenticated update on the outsourced data.

Security Analysis The modified DPDP I scheme satisfies the authenticity and freshness properties as described in Section 4 (this directly follows from the same guarantees provided by the original DPDP I). Given a probabilistic polynomial-time adversary $\mathcal{A}$ that wins the data possession game with some non-negligible probability, there exists a polynomial-time extractor algorithm $\mathcal{E}$ for the original DPDP I which can extract the challenged vectors (except with negligible probability) by interacting with $\mathcal{A}$. Now, the extractor algorithm $\mathcal{E}'$ for the modified DPDP I challenges the adversary with two different challenge sets $Q = \{(i, \nu_i)\}_{i \in I}$ and $Q' = \{(i, \nu'_i)\}_{i \in I}$ on the same commitment r, where each ν_i (or ν'_i) is a random value. Then, $\mathcal{E}'$ gets two responses of the form $B' = \sum_{i \in I} \nu_i b_i + r$ and $B'' = \sum_{i \in I} \nu'_i b_i + r$, and the extractor now forms another $B''' = \sum_{i \in I} \nu''_i b_i$ where $\nu''_i = \nu_i - \nu'_i$ for each $i \in I$. We note that $B''' = \sum_{i \in I} \nu''_i b_i$ is similar to a response from the adversary in the original DPDP I scheme described in Appendix B.1. Thus, $\mathcal{E}'$ can extract (at least) the challenged vectors in a similar fashion as done by $\mathcal{E}$.

Privacy-Preserving Audits We observe that the TPA does not have an access to the value of B. To get the value of B, the TPA has to solve either B from $g^B \bmod N$, or r from $R = g^r \bmod N$, both of which are infeasible for any probabilistic polynomial-time adversary $\mathcal{A}$, except with some probability negligible in λ. Thus, the audits are privacy-preserving in this modified scheme.